



The Perceptual Cues that Reshape Expert Reasoning

Michael Harré¹, Terry Bossomaier² & Allan Snyder³

¹The School of Psychology, The University of Sydney, Sydney, Australia, ²The Centre for Research in Complex Systems, Charles Sturt University, Bathurst, Australia, ³Department of Physiology, The University of Sydney, Sydney, Australia.

SUBJECT AREAS:

ANIMAL BEHAVIOUR

BEHAVIOUR

COMPUTATIONAL BIOLOGY

PHENOMENA AND PROCESSES

Received
19 April 2012

Accepted
18 June 2012

Published
11 July 2012

Correspondence and requests for materials should be addressed to M.S.H. (michael.harre@sydney.edu.au)

The earliest stages in our perception of the world have a subtle but powerful influence on later thought processes; they provide the contextual cues within which our thoughts are framed and they adapt to many different environments throughout our lives. Understanding the changes in these cues is crucial to understanding how our perceptual ability develops, but these changes are often difficult to quantify in sufficiently complex tasks where objective measures of development are available. Here we simulate perceptual learning using neural networks and demonstrate fundamental changes in these cues as a function of skill. These cues are cognitively grouped together to form perceptual templates that enable rapid ‘whole scene’ categorisation of complex stimuli. Such categories reduce the computational load on our capacity limited thought processes, they inform our higher cognitive processes and they suggest a framework of perceptual pre-processing that captures the central role of perception in expertise.

When chess grandmasters glance at a game they simply ‘get it’, not only do they choose better moves than lesser players but often these moves occur to them within seconds of first looking at a game¹, long before they have an opportunity for detailed search and analysis. How are they able to do this? Research on expertise highlights several key aspects. In games like chess, a high IQ is not necessary² but at least 10,000 hours³ of training is vital. Over this time 300,000 or more *chunks*⁴, small frequently occurring patterns, will be learned. This learning process will be non-linear: there will be times when skill plateaus⁵ and sharp transitional points need to be negotiated⁶. But learning chunks and coupling them with moves is not enough for good decisions. In the game of Go for example the best move predictor uses chunks and matches an expert’s choice 34% of the time⁷, insufficiently accurate for expert play by itself. To address the issue, amongst many others, of integrating local knowledge such as chunks into a global relational context the *Template Theory*⁸ of expertise was developed which models how chunks can be combined to form larger cognitive representations of the task space.

The Template Theory is a direct result of earlier work by Simon and colleagues who had considered the role of perception⁹ in problem solving, particularly the first seconds of considering a complex problem¹⁰. Template Theory addresses the primacy of perception and pattern recognition in tasks that previously had been thought to be solely the domain of logical reasoning such as search, planning and evaluation, i.e. the domain of conscious thought processes. Such conscious reasoning is characterised as slow, serial and capacity constrained whereas the perceptual processes Simon considered are fast, parallel and unconstrained in capacity¹¹. Recent work in this area has shown that *unconscious* perceptual learning can occur in domains as complex as board games¹², speech¹³ and mathematics¹⁴. In such cases early perceptual processes can adapt and learn the complex and often noisy relationships between visual elements, effectively acting as a pre-processing step that influences the later stages of cognition. Most recently this has developed into the perceptual learning of human expertise¹⁵ and is characterised by the developmental changes induced in early sensory regions of the brain by extensive experience. Such early stage adaptations change the way in which a perceiver extracts information from the environment and it is often implicit in two distinct ways: *perceptual learning* is implicit in that it is not a declarative learning process, instead it occurs without the perceiver being aware of what is being learned^{16,17}, and *perceptual expertise* is implemented without awareness in so far as the perceiver is not overtly aware of the influence their acquired knowledge has on the decisions they make^{18,19}.

Simon summarised his results in the following way²⁰: “The situation has provided a cue; this cue has given the expert access to information stored in memory, and the information provides the answer [...]” whereby “[w]e are aware of the *fact* of recognition, which gives us access to our knowledge [...]; we are not aware of the *processes* that accomplish the recognition.” [original emphasis]. The goal of this work, then, is to find the *perceptual templates* that amateurs and professionals have acquired through perceptual learning and that they implement as the basis of their perceptual expertise when playing the oriental game of Go. Such templates are reduced representations of



the state of a game, they contain a subset of the total number of pieces on the board but this subset makes up the perceptually learnable relationships in the game. The individual cues that make up a template, i.e. the positions and colours of the game pieces, are processed in parallel during the early stages of sensory perception, much as the global relationships between elements in natural scenes are processed in parallel in early stages of perception²¹. It is at this level of cognition at which templates are employed giving rise to *expert intuition*²². A principal difference between the amateurs and the professionals lies in their perception of the global context of the game in which individual moves are made^{23,24}. In discussing the role of these templates we will use recent results on early perceptual learning and the pre-processing of visual stimuli to suggest that experts use the rapid recognition of complex patterns, mediated by perceptual templates, in order to efficiently constrain and guide their search for good moves.

Visual illusions highlight the subtle and persistent nature of such perceptual pre-processing. The Ponzo illusion shown in figure 1 comes about through very early processing in the visual cortex, area V1, and the subjective impact of the illusion is influenced by the surface area of an individual's V1 region²⁵, highlighting the early stages at which such effects occur and how they are influenced by gross neural properties. In figure 1a. the illusion is that the top red bar is longer than the bottom red bar and it is induced by the parallel railway lines that appear to draw closer together in the distance. The key is that the two red bars appear to be placed at different distances from the observer, a perspective strongly informed by the relationship between the converging tracks and the red bars. The converging tracks act as cues that inform the observer of the different scales of objects in different parts of the scene, and so the apparent differences in the size of the bars is coherent with respect to these cues. Figure 1b. shows how a sparse representation retaining only the two contextual cues (the two lines that converge) and the relevant information (the two red bars being compared) is able to maintain the sense of the illusion when no other information is retained. Even when we are consciously aware of being deceived by the illusion, such overt awareness does not easily change the sense of the illusion, demonstrating how the early stages of perception that use such learned cues are not

readily switched off. But when the contextual cues are removed (figure 1c.) the illusion vanishes: the context informs us of a particular interpretation of the environment and while this guidance is often a very useful heuristic at times it can lead us astray.

Such illusions are most likely at least partly a result of early cognitive processes reducing the vast amount of information we receive from the environment¹⁹. It has long been recognised that the information capacity of short term memory is tightly constrained²⁶ and it is now thought to hold only a few elements²⁷. This bottleneck, called the Working Memory (WM)^{28,29}, sets an upper limit on how much information can be held in active memory at any given time. WM capacity is not thought to be improved by task specific learning such as chess training³⁰ or by perceptual expertise³¹, although generic (i.e. non-task specific) training techniques have been shown to improve WM capacity³². In order to manage these limitations early perceptual processing does not expand the capacity of WM, instead it reduces the amount of information being passed to our WM, capturing only the relevant information necessary for higher order processing. In the Ponzo illusion of figure 1 the environment is reduced to the commonly occurring regularities that would usually inform our scene comprehension and set the context, i.e. the two converging lines, as well as the information relevant to the task, the parallel red bars. In this sense the contextual information can be thought of as a general purpose caricature of the scene that can be applied flexibly in many different circumstances in which the omitted details are not immediately relevant.

Such considerations suggest an important difficulty our cognitive processes are able to address: the total information contained in an environment is too vast to process using deliberative reasoning but the information contained in localised chunks of the environment is too focused to be useful by itself. What we would like to find are the small number of visual cues that make up the salient aspects of the environment and show how these cues change as a function of skill. One approach is to use a neural network that can automatically form ordered, compressed representations of sensory perceptions such as Self-Organizing Maps (SoMs)³³. SoMs have been used as a model of neurological organisation³⁴ as well as a tool for data-mining³⁵ and




Figure 1 | The original Ponzo illusion is shown in 1.a. The corresponding Perceptual Template (the converging lines) is shown in 1.b along with the red bars about which a decision regarding their relative lengths needs to be made. In 1.b the detailed information has been removed but the information that informs our judgement of the length of the red bars is retained. Without the perceptual template the illusion vanishes as can be seen in 1.c.



they have the benefit of being unsupervised learners³⁶, that is to say they extract structural regularity from data without external guidance. This last point is significant, from a behavioural perspective human players implicitly learning relationships between game pieces are not aware of what is being learned, they are only picking up on the statistical regularities in the task environment¹⁷, this requires an unsupervised process.

An SoM can be thought of as a non-linear extension of principal component analysis³⁷ based upon biological principles³⁴. The key idea is that a complex perception of the environment containing many different elements, encoded as a vector $\mathbf{x}^i = [x_1, x_2, \dots, x_n]$, is processed in parallel by a large array of neurons, as previously theorised³⁸ and recently observed in the laboratory³⁹, whereby each idealised neuron represents a learned model of the environment. Algorithmically, an SoM neuron (a model \mathbf{m}^j) is a vector of the same dimensions as the vectors that represent a perception of the environment: $\mathbf{m}^j = [m_1, m_2, \dots, m_n]$ for $j \in \{1, \dots, N\}$ neurons in the SoM. When an SoM is presented with a training vector \mathbf{x}^i of the environment it is compared to all SoM neurons and one, the winning neuron \mathbf{m}^c , is that which satisfies:

$$\delta(\mathbf{x}^i, \mathbf{m}^c) = \min_j(\delta(\mathbf{x}^i, \mathbf{m}^j)) \quad (1)$$

where $\delta(\mathbf{v}^1, \mathbf{v}^2)$ is usually the Euclidean distance between vectors \mathbf{v}^1 and \mathbf{v}^2 . This is the metric used in the SoM MATLAB toolbox developed

at the Helsinki University of Technology⁴⁰ and it was their implementation used in this work.

Initially all of the neurons \mathbf{m}^j have randomised elements. After selecting \mathbf{m}^c the weights of neurons $\mathbf{m}^k, k \in \{\text{index of neurons 'local' to and including neuron } c\}$, are updated:

$$\mathbf{m}_{\text{after}}^k = \mathbf{m}_{\text{before}}^k + \kappa \cdot \alpha \cdot \delta(\mathbf{x}^i, \mathbf{m}^k) \quad (2)$$

The functions κ and α refer to the convolution kernel and the learning rate³³ respectively and while they have a temporal dependency within the SoM toolbox they were otherwise left fixed throughout this work. This algorithm updates the best matching neuron in the SoM and all neurons locally connected to it, see figure 2 and the Methods section for a complete description of the SoM implementation.

Our task environment is the choice of moves made in the oriental game of Go. Go is played on a board made of a 19×19 grid on the vertices of which game pieces, called *stones*, are placed. Before the game starts, players choose to play with either black or white stones and they take it in turns to place a stone on one of the vertices of the board. The goal is to capture more territory than your opponent by surrounding regions of the board with stones of your own colour that are connected to each other in chains. A chain is formed by stones that are placed directly next to each other in the north, south, east or west direction on the grid. Stones are only removed when they are completely surrounded by the stones of the other player such that

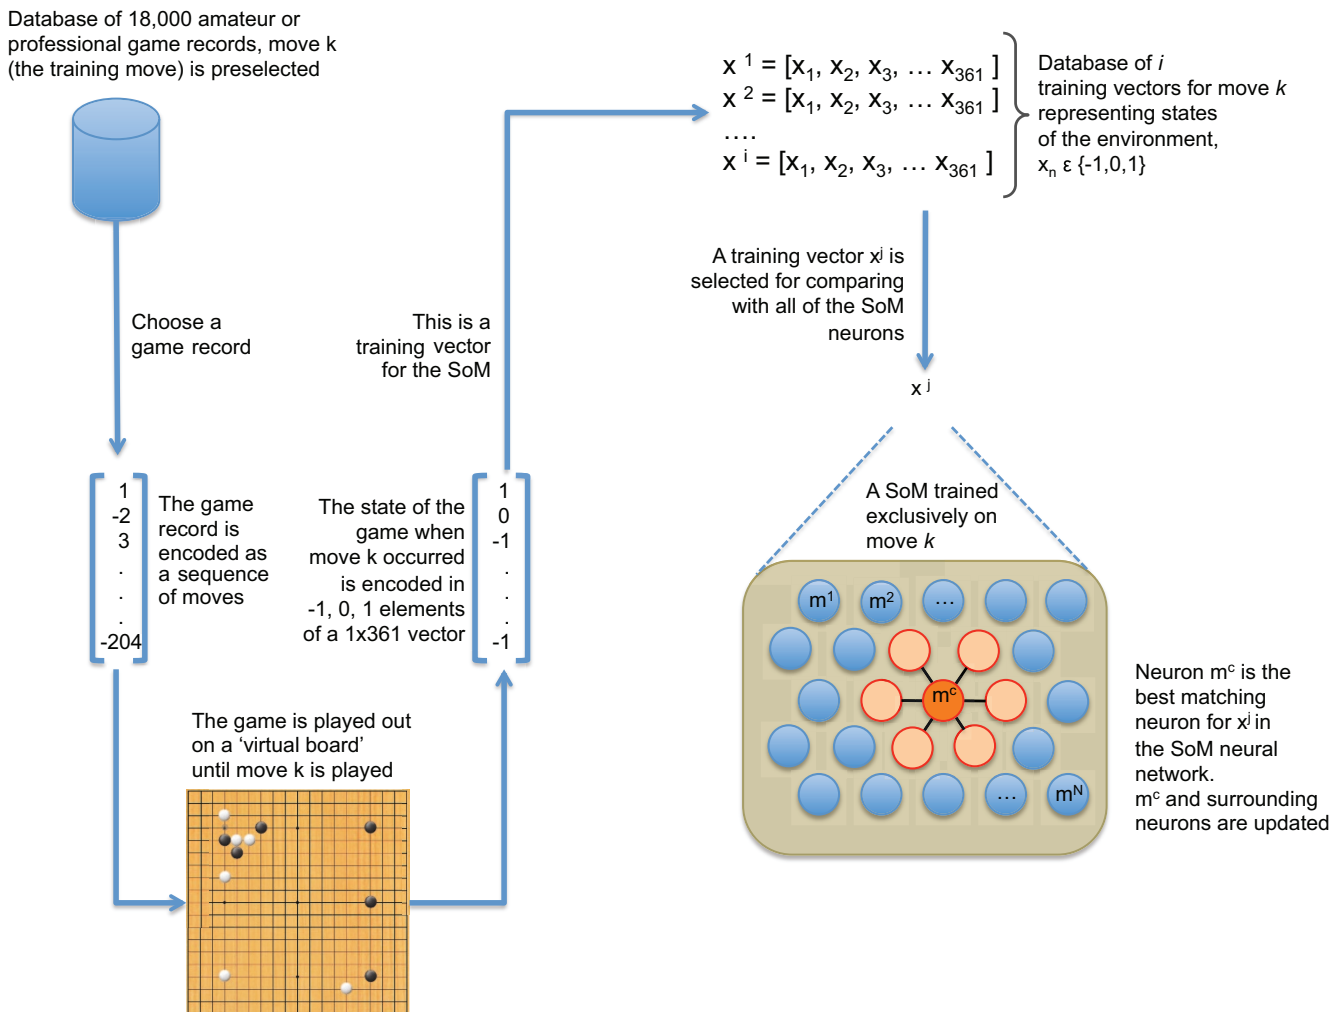


Figure 2 | The computational implementation of Self-Organizing Maps for extracting the structured information from Go games given a move k . The resultant trained SoM output at the end contains neurons with continuous weights, it is a further step to threshold these internal neuronal weights in order to generate the unique templates.



there are no more free vertices within the surrounded territory for the surrounded player to play on. At the end of the game the player with the most territory surrounded wins.

At any given point in the game there is a maximum of 361 possible vertices on which to place a stone, or alternatively there are a maximum of 360 positions that can influence any given move. In order to reduce this number we need to find those stones that frequently co-occur when a given move is made, these stones make up the contextual cues and a combination of these cues is an instance of a perceptual template. Note that for any move there will be many different board configurations in which it occurs, and as a template is a reduced representation of the state of the board (containing as it does only those stones that have occurred frequently together) then any given template can fit multiple game instances. In this sense SoMs³³ generate perceptual templates that categorise the different ‘game scenes’ in which moves are made. This constancy in the relationship between the multiple cues and the move itself is necessary for an expert’s perceptual learning to occur⁴¹ and it is the basis on which our SoMs are able to categorise high dimensional data⁴², also see the Discussion below.

Each SoM neuron is a 1×361 vector representing a ‘model’ of the Go board when a move was made, each element within a neuron is a learned weight w_i in the interval $[-1, 1]$ representing how ‘black’ or ‘white’ each board position is in that neuron’s model. Instead of using these continuous values we set a cut-off value using a threshold parameter t : $|w_i| > t$. This cut-off restricts the elements of each neuron to the discrete values $\{-1, 0, 1\}$ so each neuron encodes a model of the game containing only black, white or empty positions based on the threshold. The unique set of these neurons are the perceptual templates⁴², they represent a collection of different contextual models of the game environment and are the reduced representations of the state of the game corresponding to the structured regularities that player’s are repeatedly exposed to each time they make a move during game-play.

Results

Table 1 shows the template statistics: the total count of templates for amateurs and professionals, the percentage increase in the number of templates that a professional has compared to amateurs, the percentage of templates that are shared between amateurs and professionals and the size (in terms of number of stones) of the largest templates. Most notably the difference in the number of templates for amateurs and professionals is quite small ranging from around 12% to 38%. However the number of templates that are common to both amateurs and professionals is equally small, ranging from around 12% to 19%. There is also a persistent difference in the maximum size of amateur and professional templates, professionals being more than 100% larger at some threshold values although there are relatively few of the largest templates for either amateurs or professionals (see figure 3).

The distribution of the number of stones in each template for four different threshold values is plotted in figure 3. For both amateurs and professionals the mean number of stones was relatively small, between approximately 3 and 6 stones and was more insensitive to changes in the threshold parameter than might be expected. On the other hand the tails for these distributions are quite different for the two classes of players. Table 1 shows the maximum size differs

greatly, but this is caused by a very small number of larger templates. A closer look at the templates showed that the professionals used a small number of localised patterns of stones so frequently that they were learned by the SoMs, something that happened considerably less often for the amateurs. However the central portion of the probability distributions in figure 3 remains qualitatively very similar across threshold values and player class.

Figure 4 measures the intersection between the amateur and professional templates as a function of the professional template indexing (the indexing is discussed in the Methods section). As the index increases (i.e. the size and complexity of the templates increases), the rate of change in the size of the intersection set decreases until eventually adding a new professional template does not increase the size of the intersecting set. A gradient of 1 in this curve implies that for each professional template added there is a corresponding amateur template, near the origin a sustained gradient of 1 is clear but as the indexing increases the gradient progressively decreases, indicating that the more complex a professional template is the less likely it will also be in the amateur set of templates.

Discussion

The goal of this work has been to find and compare the structured information, in the form of contextual cues, that is available to experts and non-experts in the game of Go. It is argued that this information is used during implicit learning and subsequent early perceptual processing of information within a given domain of expertise to aid in fast and accurate categorisation and decision-making in complex environments. In particular, these processes enable the reduction of the dense information perceived in a complex natural environment using the available structured regularities. Furthermore, the integration of these cues into a cognitive whole leads to the notion of perceptual templates, the aggregate, sparse representations of the salient features of the task environment that enables many of the remarkable feats reported in studies of domain-specific expertise.

Figure 5 shows a cognitive model that demonstrates how such perceptual templates might be implemented. The entire scene is initially processed by low level visual systems⁴³ combined with perceptual templates to produce a perceptual whole in a very short period of time. Estimates of the length of time it takes to categorise the ‘gist’ of complex scenes range from about 30 milliseconds up to around 150 milliseconds^{44–47}. Note that in the study by Thorpe et al.⁴⁴ presentation times of the images (20 ms) were too brief to allow eye movements to search the image, effectively requiring the subjects to comprehend a complex image as a perceptual whole. This initial ‘feed forward sweep’⁴⁵ of perceptual information is too quick for neural feedback pathways to influence the scene perception, suggesting that strictly feedforward¹⁸ processes of complex visual scenes are sufficient for early perceptual categorisations. Recent work on the physiological basis of expertise, both theoretical⁴⁸ and empirical⁴⁹, provides support for cognitive templates being located in the inferior temporal cortex. It is this region that fMRI⁵⁰ and diffusion tensor imaging⁵¹ studies have strongly implicated the visual perception of Go board patterns in experts but not novices.

In this sense our ability to form a categorical impression of a complex scene is almost immediate and it is this categorical impression that perceptual templates capture. The implementation process

Table 1 | Threshold dependent properties of templates

Threshold	Am. : Prof. count	Intersection count	% increase Am. to Prof.	% of shared templates	Am.: Prof. max. size
0.95	3,353 : 4,618	889	37.7%	19.3%	11 : 24
0.90	8,413 : 10,930	1,680	29.9%	15.4%	20 : 30
0.85	15,215 : 18,185	2,399	19.5%	13.2%	35 : 82
0.80	23,458 : 26,319	3,119	12.2%	11.8%	56 : 109

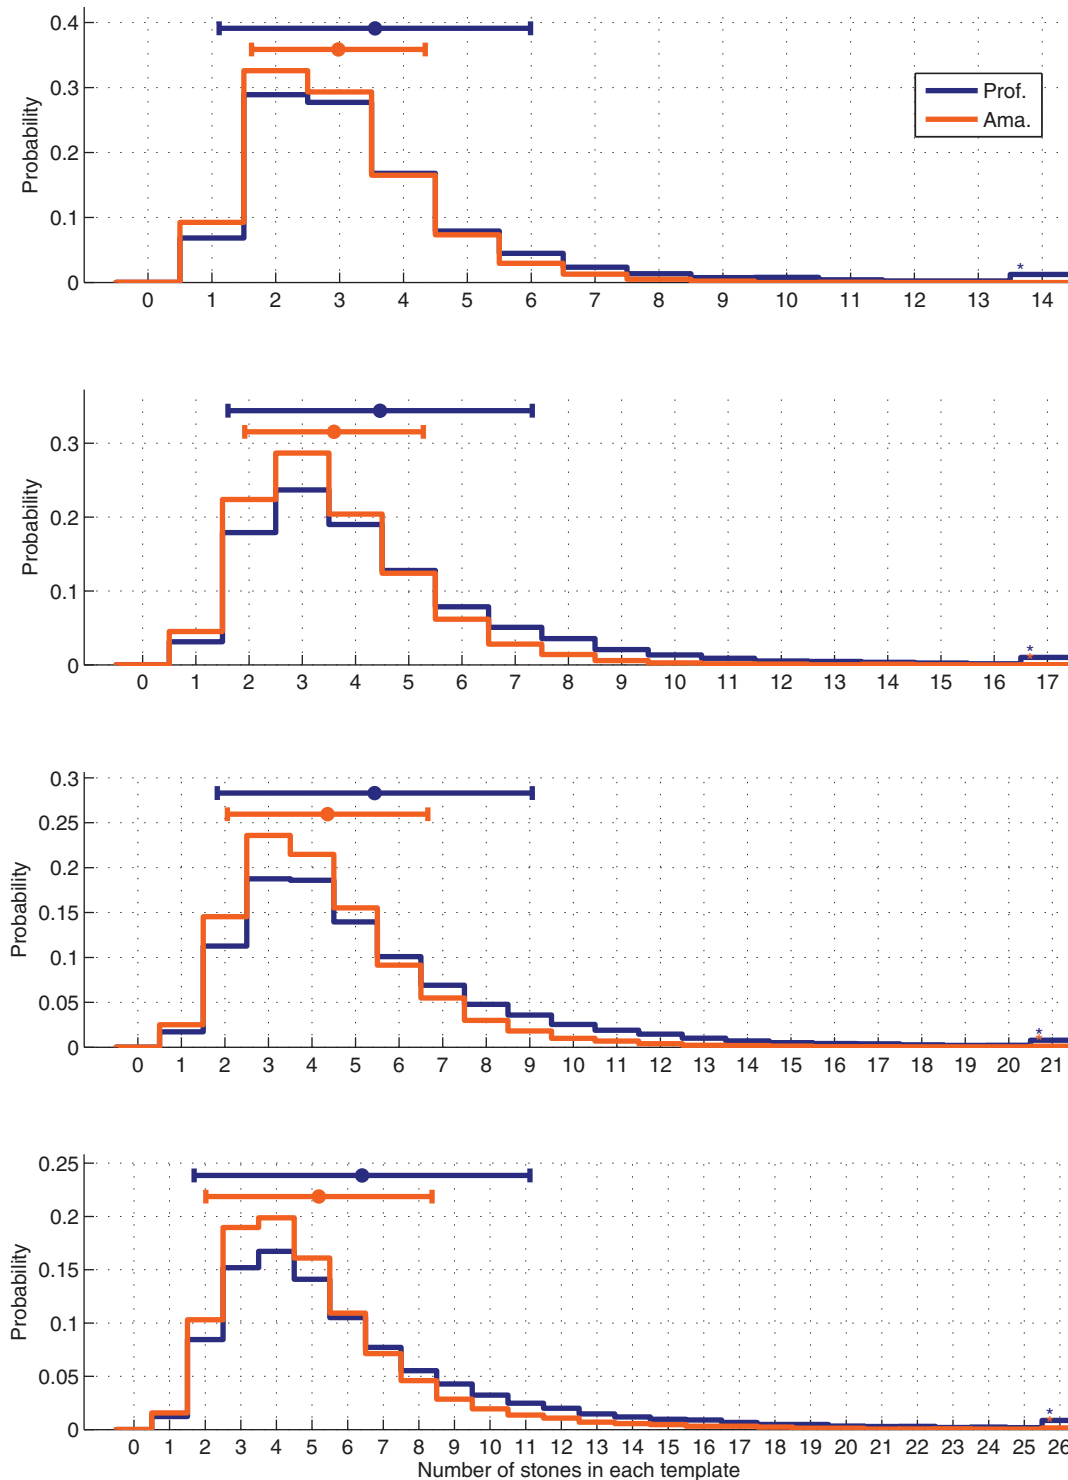


Figure 3 | The probability distributions for the sizes of the templates (the number of stones each template contains). From top to bottom the threshold values are: 0.95, 0.90, 0.85 and 0.80. The mean and standard deviation of each distribution are plotted on interval lines above each plot.

is as follows: In figure 5 the combination of the four cues A, B, C and D are compared in parallel to all of the perceptual templates (simplified models of the world) the perceiver has learned and template 3 minimises the difference between the model it encodes and the current visual environment and templates 1 and 2, for example, do not. This results in an activation of a single template (template 3) that acts to contextually activate a later network of processing modules, i.e. which of the modules x , y or z should be activated to provide for further analysis.

Template 3 activates the initial eye saccade to some region of the scene (e.g. to cue B) for more deliberate processing in a serial fashion. A combination of contextual information and localised analysis, based on higher level cognitive outputs, may lead to further eye saccades that allow for greater analysis of the environment. Such detailed analysis is usually an evaluative process requiring a small number of alternative strategies to be maintained in working memory at the same time. In this sense there is considerable conceptual similarity between this model and that of both CHREST⁵² and the

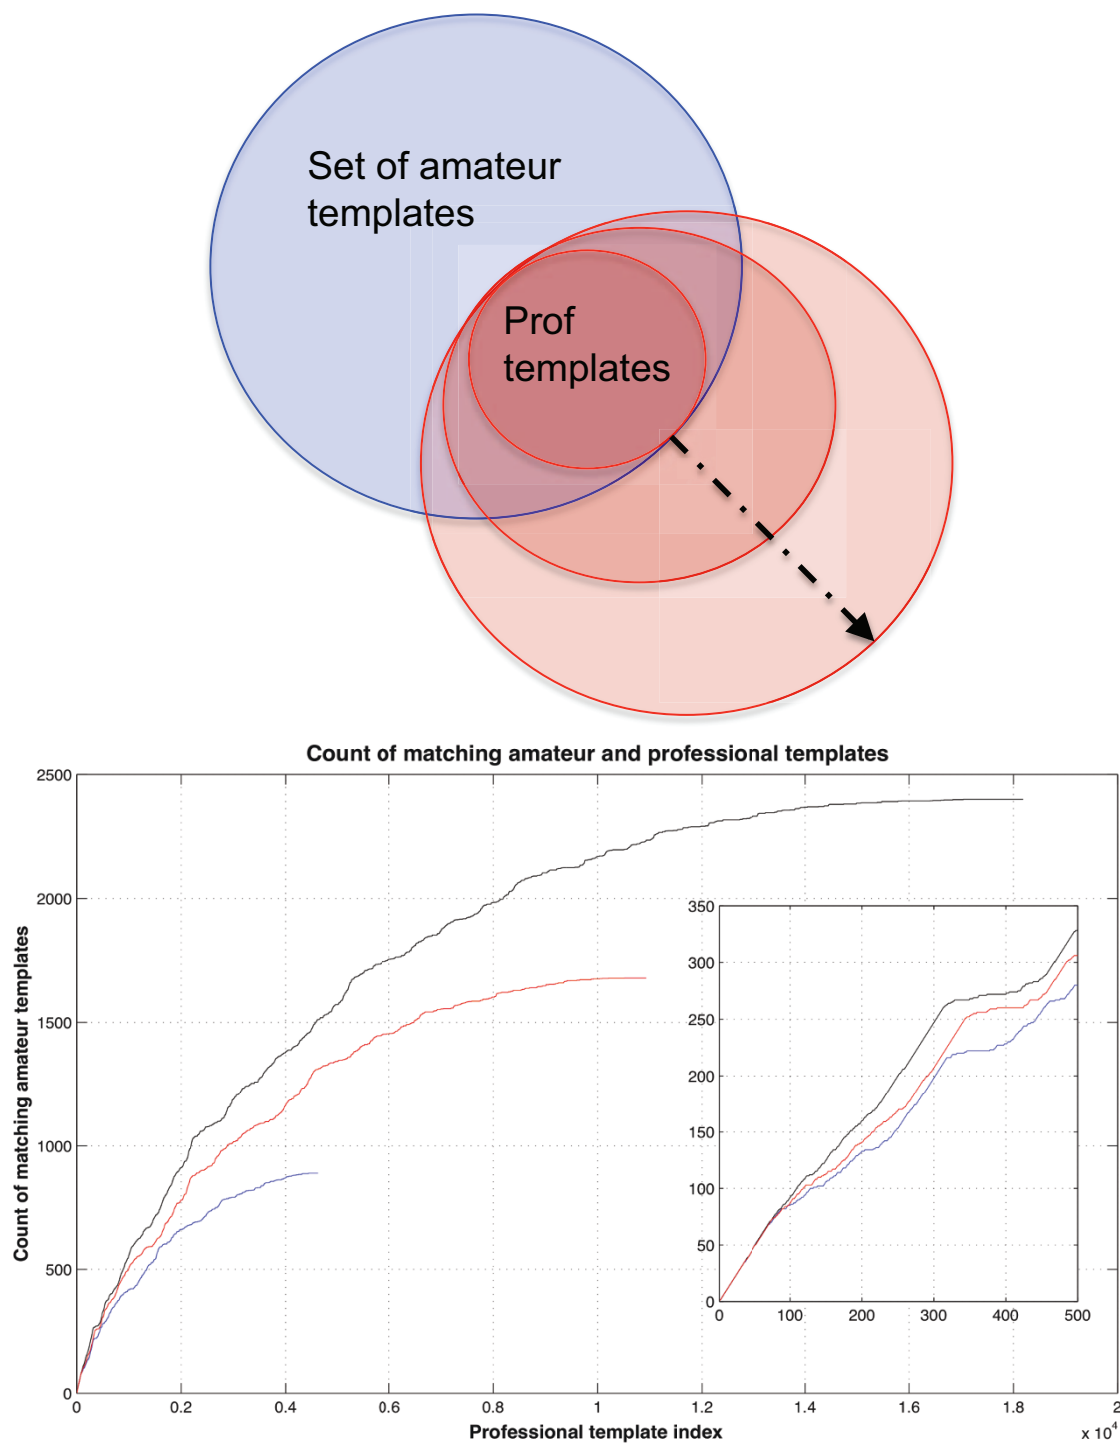


Figure 4 | The intersection between experts and non-experts. Top: A schematic representation of the intersection between sets of templates. The set of amateur templates are held constant while the set of professional templates is allowed to increase. The size of the professional set is increasing in the positive direction of the x -axis in the bottom plot, i.e. the larger the professional set the larger and more complex the templates are that have been included. Bottom: The size of the intersecting set of amateur and professional templates as a function of the professionals' indexing. Three different thresholds shown, from bottom to top: 0.95, 0.90 and 0.85. Inset: An expansion near the origin.

guided search³⁸ models. Note that changes in the perceptual templates will result in changes in the patterns of the eye saccades that are related to the development of expertise^{53,54}. This early processing of the context persistently influences perception, just as visual illusions do, and provides the necessary categorical information required to constrain later search heuristics and the evaluation of moves in order to keep the computational load of such tasks within the bounds of our limited cognitive capacities.

In light of the earlier discussion of visual illusions, it is known that an illusion's effect decreases over the time course of perception, illusions being strongest in the first stages of perception and then modulated by later, higher order, cognitive processes⁵⁵. On the basis of this evidence and that of the role of V1 on an illusion's subjective impact²⁵ and contextual relationships in visual scenes⁵⁶, it is reasonable to suggest that the categorisation of a scene happens relatively early in visual processing and is modulated by later, top-down processes

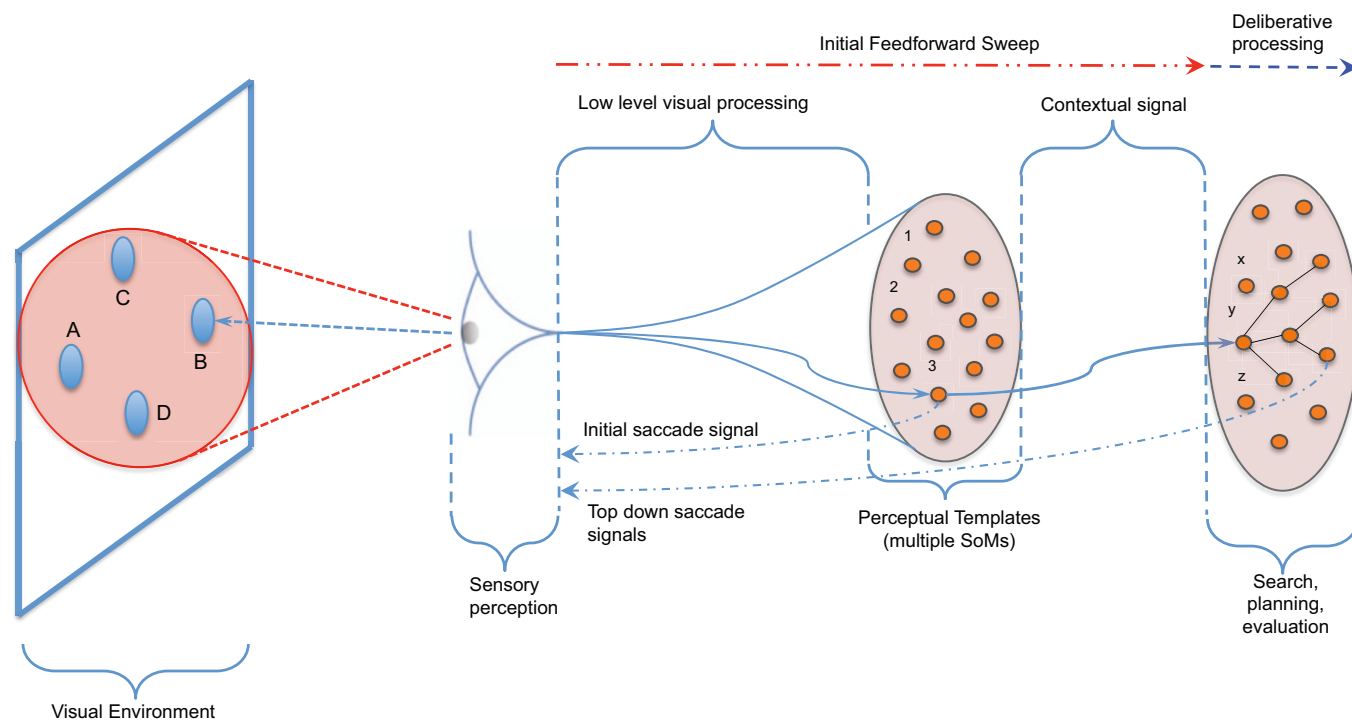


Figure 5 | A cognitive model of the implementation of perceptual templates. Early processing encodes the environment such that it can be compared in parallel with a large number of neural representations of templates. A single template is found that is a 'best fit' given the current environment. This template signals where the eye's gaze should saccade to first as well as providing the context that informs later cognitive processing. Amateurs and professionals differ in their perceptual templates resulting in different interpretations of the context and different patterns of eye saccades.

that enable a more precise comprehension of the local characteristics of the scene. This is very similar to the scene-centred approach to understanding the holistic properties of a scene, called the 'gist', recently put forward by Oliva and colleagues^{21,43,57}. There is a significant difference though in that Oliva et al's work is based upon natural scene analysis and not strategic games, but further work is expected to clarify the similarities in these different approaches. The fact that the same mechanisms that are in play in game expertise might also be in play during natural scene comprehension is an exciting possibility that suggests a very general mechanism may mediate an exceptionally broad range of complex task environments.

As mentioned above, there is considerable conceptual similarity between our results on perceptual templates and the research of Gobet and others on Template Theory, however the two are not synonymous and some important distinctions should be made. Template Theory developed out of chunking theory as a theoretical construct by Chase and Simon⁵³ and was based on the earlier work of de Groot⁵⁸ on chess expertise and Miller²⁶ on capacity limits in our cognitive processes. In the original Template Theory⁸, chunks containing several chess pieces are learned by novices but as their experience grows so too does the size, in terms of the number of pieces they contain, of these chunks. Chunks are stored in long term memory but pointers to these chunks are held in short term memory that can only hold around 3 such pointers due to capacity restrictions. As the chunks grow in size the number of pointers does not increase but the size of the chunks they point to do, thereby allowing experts access to greater amounts of information and circumventing the limited capacity of our working memory. Templates are larger and more elaborate structures than chunks, they contain 15 to 20 game pieces⁴ but they also have slots into which smaller chunks can be inserted⁸. A template then is an example of a "schema" as studied in psychology where they are "... implicitly learned in the process of acquiring substantive knowledge ..."⁴

Much of the high level description of Template Theory is similar to the perceptual templates of this study: perceptual templates contain a

reduced number of game pieces (the *core* in Template Theory) that are implicitly learned during the course of acquiring expertise, they are augmented by detailed and localised analysis of the board (similar to the role chunks have with respect to slots), they are composed of consistently co-occurring game pieces that augment strategising, move selection and circumvent some of our cognitive limitations. The most significant difference between the two lies in the method of extracting the templates. The CHREST cognitive architecture that implements Template Theory uses a 'roving eye' to scan many chess games in order to build chunks first and then more elaborated structures that eventually become templates⁵⁹. That is to say that Template Theory builds up from chunks to form more elaborated structures containing a core and slots that can then contain a variety of different chunks. There is considerable empirical support for this model⁵⁹.

On the other hand the cognitive implementation of perceptual templates acts much more like 'SoM-filtered' Bayesian inference. It does not start from chunks and build up like Template Theory, instead it takes the whole board as a single perceptual input. Each training vector x^i is a whole board configuration from which a move was then made: $x^i \rightarrow m_k$ where the training vector x^i varies from instance to instance but the move m_k does not, see figure 2. This implicitly conditions the training vector on the move that was then made. Given thousands of training vectors conditioned on a fixed move a SoM (a single 50×50 neural network) learns to categorise the board configurations according to the frequently occurring game pieces, filtering out all the infrequently occurring pieces.

This implementation is quite different from Template Theory, it implies that *when a certain move is made*, we implicitly learn the statistical regularities associated with that move. The resultant templates can then be used to invert this process: when a board configuration is perceived, for example during a game, and early perceptual processes are required to suggest a few possible moves (as well as communicating contextual-categorical information) the templates compete amongst themselves, most likely based on a competitive activation model⁶⁰, to communicate to higher cognitive processes



all of the moves that had learned that particular template. This is because while the mapping $x^i \rightarrow m_k$ fixes the move during one SoM's learning, there are other SoMs trained on different moves using different training vectors that may have learned the same perceptual template, i.e. a single template might have been learned by multiple moves resulting in multiple possible next moves being generated from one template. Given the considerable differences in these two different template paradigms it is not clear where the similarities and the differences between the resultant templates lie.

There is already some interesting evidence suggesting a difference in the two methods. The largest templates found using the highest threshold parameter ($t = 0.95$) was 24 stones (see table 1). By comparison Gobet and Simon reported psychological experiments⁴ showing that chess Masters have a maximum of around 15 chess pieces in their chunks/templates. This most likely occurs because chess positions are less stable than those of Go, stones in Go remain in place unless captured which happens rarely when compared with how often chess pieces move. This means that larger templates can be learned more readily in Go because of the perceptual regularity of the game pieces. A useful exercise for further study will be to implement a SoM based perceptual templates analysis of chess positions.

Furthermore, future research into the role of such templates in expert cognition should also be critically informed by psychological experiments. For instance board configurations that more closely match professional templates should result in more rapid generation of possible next moves, i.e. perceptual templates should increase the fluency of move generation for experts. Similarly, board configurations that do not easily match any perceptual templates should increase the time it takes a player to generate options for the next move. Such experiments will help establish the psychological validity of perceptual templates and further inform their theoretical development.

These results paint an intriguing picture of the perceptual templates available to skilled and unskilled practitioners. While there are many thousands of unique templates ('game scene' categories), this still represents a massive reduction in the total number of possible scenes that would otherwise need to be analysed deliberately at all levels of detail, much as many artificial intelligence systems do. However, despite the striking similarities in several high-level properties, such as total number (table 1) and distributions of sizes (figure 3), the overlap between amateur and professional players is small and they systematically diverge for the larger and more complex templates (figure 4). From this we see that a professional's perceptual learning in the game of Go is informed by quite different information to that of an amateur and that they share only the most basic information. This is not a sufficient explanation for all of expertise: these templates provide an approximate analysis of the game, they still need to be connected with later cognitive processes and ultimately with a decision regarding where to move. In this light, the current work provides novel evidence of a measurable mechanism for some of the remarkable differences in performance between expert and non-expert decision-making in complex tasks.

Methods

Game data and preprocessing. We used 18,000 games of professional ranked players (rank 5 dan professional and above) and 18,000 games of amateur players (rank 1 kyu, 1 dan or 2 dan amateur). The professional games were part of the GoGod database available commercially at www.gogod.co.uk and the amateur games were recorded during online play from the KGS Go server: www.gokgs.com.

Each game was converted into a $3 \times m$ matrix where m denotes the number of moves played during the game, each $m_i = [x, y, \pm i]$ where $x, y \in \{1 \dots 19\}$ are board co-ordinates and $\pm i$ is the move number (a negative i represents a black move on the i^{th} turn, a positive i represents a white move). This is sufficient to encode an entire game as a sequence of moves, but it ignores possible captures where previously played stones are removed from the board, freeing up positions that can be played later in the game. While this does not affect the encoding of the game (some positions might be played more than once during a game, but this is irrelevant for the sequence of moves played), it does have an effect on the learning vectors that are presented to the SoM

where stones have been removed after capture. This issue is addressed at the point at which the game is 'played out' during the SoM analysis discussed below.

SoM implementation. Figure 2 provides a diagrammatic representation of the implementation. In order for an SoM to learn which stones are commonly present when a move is made, the state of the game when that move was made needs to be encoded in a 1×361 vector (a linearised representation of the 19×19 board) where each element of this vector was $+1$, -1 or 0 representing a white stone, a black stone or an empty position, respectively.

Starting with either the professional or the amateur database of games, we first nominate a (linearised) position on the board, position $k \in \{1, \dots, 361\}$. A game record is then chosen from the database and the game is recorded in the sequence in which the moves were played in a 19×19 matrix (representing the current state of the board, initially all entries set to 0) where either $+1$ or -1 is recorded depending on whether a move was white or black, respectively. Each new move is checked to see if it is a capturing move, if so all of the corresponding stones that are captured are removed from the matrix and the game continues. This is repeated until a move is made at position k (but is not yet recorded in the matrix). If the move at position k is a white move, the state of the game is left unchanged, if it is black the game-state is multiplied by -1 , essentially making all moves at position k a 'white' move. This does not change the strategic relationships between the stones but it does prevent the SoM from learning separate templates for white and black moves. The game is then stopped and the current (linearised) state of the board is then the training vector for the SoM. In practice this means the training vectors are 1×361 vectors containing ± 1 and 0 elements representing the state of the game when a move was made at position k . Note that each initial game record has a length equal to the number of moves and so changes from one game to the next. However the training vectors representing a board configuration are all of the same length: 1×361 , enabling them to be compared with the 'world model' encoded by each SoM neuron.

This procedure requires each of the 18,000 amateur and professional games to be played until move k is found, but in some games k is never played in which case the game record is not used leading to slightly fewer training vectors. Each training vector is an input into a 50×50 neuron SoM that is dedicated to learning the board patterns when move k is made. This procedure is repeated for all $k \in \{1, \dots, 361\}$, resulting in an aggregate SoM neural network containing $50 \times 50 \times 361 = 902,500$ neurons, where each neuron is a 1×361 vector representing the learned real valued weights in the interval $[-1, 1]$ for each board position.

There are two of these aggregate SoM networks, one for the professionals and one for the amateurs. These networks were too large to analyse directly so the learned weights were set a threshold value t (described below), different values of which were used to generate the results and this also significantly reduced the size of the datasets we had to analyse, see Table 1.

Thresholding and sorting templates. In order to see where the most significant differences in the templates lie they were sorted and indexed in three steps. A MATLAB script takes a list of templates and first finds the unique templates (i.e. after thresholding of the learned weights, the built-in MATLAB function `unique []` removes repeated templates and sorts them in ascending order), then by frequency of each stone's occurrence in list and finally by the number of stones in each template. In the following script, list contains 902,500 vectors (trained SoM neurons) of size 361×1 with real-valued elements in the interval $[-1, 1]$ and t is the threshold value. The list that is output in the final step is a reduced set of unique, sorted templates with discrete elements containing values $\{-1, 0, 1\}$ representing the position and colour of stones on a linearised board:

```
list (list > t) = 1;
list (list < -t) = -1;
list (list >= -t & list <= t) = 0;
list = unique [list];
[~, index] = sort [sum [abs [list], 1], 'descend'];
list = sortrows (list, -index);
[~, index] = sort [sum [abs [list], 2], 'ascend'];
list = list (index, :);
```

The significant steps are template size (templates with fewer stones have lower index values) and overall frequency of occurrence (lower index values are given to templates in which their component stones occurred most often).

- Chase, W. & Simon, H. Skill in chess. *Am Sci* **61**, 394–403 (1973).
- Ericsson, K. & Lehmann, A. Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Ann Rev Psychol* **47**, 273–305 (1996).
- Ericsson, K., Krampe, R. & Tesch-Römer, C. The role of deliberate practice in the acquisition of expert performance. *Psychol Rev* **100**, 363 (1993).
- Gobet, F. & Simon, H. Five seconds or sixty? Presentation time in expert memory. *Cognitive Sci* **24**, 651–682 (2000).
- Ericsson, K. The scientific study of expert levels of performance: general implications for optimal learning and creativity. *High Abil Stud* **9**, 75–100 (1998).
- Harré, M., Bossomaier, T. & Snyder, A. The development of human expertise in a complex environment. *Mind Mach* 449–464 (2011).
- Stern, D., Herbrich, R. & Graepel, T. Bayesian pattern ranking for move prediction in the game of Go. In *Proc 23rd Int Conf Mach Learning*, 873–880 (ACM, 2006).
- Gobet, F. & Simon, H. Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychol* (1996).



9. Simon, H. Theories of decision-making in economics and behavioral science. *Am Econ Rev* **49**, 253–283 (1959).
10. Simon, H. & Barenfeld, M. Information-processing analysis of perceptual processes in problem solving. *Psychol Rev* **76**, 473 (1969).
11. Kahneman, D. A perspective on judgment and choice: Mapping bounded rationality. *Am Psychol* **58**, 697 (2003).
12. Wan, X. *et al.* The neural basis of intuitive best next-move generation in board game experts. *Science* **331**, 341 (2011).
13. Toscano, J. & McMurray, B. Cue integration with categories: Weighting acoustic cues in speech using unsupervised learning and distributional statistics. *Cognitive Sci* **34**, 434–464 (2010).
14. Kellman, P., Massey, C. & Son, J. Perceptual learning modules in mathematics: Enhancing students pattern recognition, structure extraction, and fluency. *Top Cogn Sci* **2**, 285–305 (2010).
15. Kellman, P. & Garrigan, P. Perceptual learning and human expertise. *Phys Life Rev* **6**, 53–84 (2009).
16. Lewicki, P., Hill, T. & Czyzewska, M. Nonconscious acquisition of information. *Am Psychol* **47**, 796 (1992).
17. Turk-Browne, N., Scholl, B., Chun, M. & Johnson, M. Neural evidence of statistical learning: efficient detection of visual regularities without awareness. *J Cognitive Neurosci* **21**, 1934–1945 (2009).
18. Koivisto, M. & Revonsuo, A. Preconscious analysis of global structure: Evidence from masked priming. *Vis Cogn* **11**, 105–127 (2004).
19. Evans, J. Dual-processing accounts of reasoning, judgment, and social cognition. *Annu Rev Psychol* **59**, 255–278 (2008).
20. Simon, H. What is an “explanation” of behavior? *Psychol Sci* **3**, 150 (1992).
21. Oliva, A. & Torralba, A. Building the gist of a scene: The role of global image features in recognition. *Prog Brain Res* **155**, 23–36 (2006).
22. Gobet, F. & Chassy, P. Expertise and intuition: A tale of three theories. *Mind Mach* **19**, 151–180 (2009).
23. Harré, M., Bossomaier, T., Gillett, A. & Snyder, A. The aggregate complexity of decisions in the game of Go. *Eur Phys J B* **80**, 555 (2011).
24. Bossomaier, T. & Harré, M. Global context influences local decisions. In *COGNITIVE 2011, The Third International Conference on Advanced Cognitive Technologies and Applications*, 25–29 (2011).
25. Schwarzkopf, D., Song, C. & Rees, G. The surface area of human V1 predicts the subjective experience of object size. *Nat Neurosci* **14**, 28–30 (2010).
26. Miller, G. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol Rev* **63**, 81 (1956).
27. Cowan, N. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behav Brain Sci* **24**, 87–114 (2001).
28. Baddeley, A. & Hitch, G. Working memory. *Psychol Learn Motiv* **8**, 47–89 (1974).
29. Baddeley, A. *et al.* Working memory: Looking back and looking forward. *Nat Rev Neurosci* **4**, 829–839 (2003).
30. Gobet, F. Expert memory: A comparison of four theories. *Cognition* **66**, 115–152 (1998).
31. Scolari, M., Vogel, E. & Awh, E. Perceptual expertise enhances the resolution but not the number of representations in working memory. *Psychon B Rev* **15**, 215–222 (2008).
32. Jaeggi, S., Buschkuhl, M., Jonides, J. & Perrig, W. Improving fluid intelligence with training on working memory. *P Natl A Sci USA* **105**, 6829 (2008).
33. Kohonen, T. Self-organized formation of topologically correct feature maps. *Biol Cybern* **43**, 59–69 (1982).
34. Kohonen, T. Physiological interpretation of the self-organizing map algorithm. *Neural Networks* **6**, 895–905 (1993).
35. Kohonen, T., Oja, E., Simula, O., Visa, A. & Kangas, J. Engineering applications of the self-organizing map. *P IEEE* **84**, 1358–1384 (1996).
36. Sejnowski, T. *Unsupervised learning: foundations of neural computation* (The MIT press, 1999).
37. Kohonen, T. & Honkela, T. Kohonen network. *Scholarpedia* **2**, 1568 (2007).
38. Cave, K. & Wolfe, J. Modeling the role of parallel processing in visual search. *Cognitive Psychol* **22**, 225–271 (1990).
39. Smith, S. & Häusser, M. Parallel processing of visual space by neighboring neurons in mouse visual cortex. *Nat Neurosci* **13**, 1144–1149 (2010).
40. Vesanto, J., Himberg, J., Alhoniemi, E., Parhankangas, J. & Oy, L. SOM toolbox for Matlab. *Techn. Ber., Helsinki University of Technology* (2000).
41. Garrigan, P. & Kellman, P. Perceptual learning depends on perceptual constancy. *P Natl Acad Sci USA* **105**, 2248 (2008).
42. Harré, M. & Snyder, A. Intuitive expertise and perceptual templates. *Mind Mach* **1–16** (2011).
43. Oliva, A. Gist of the scene. *Neurobiology of Attention* 251–256 (2005).
44. Thorpe, S., Fize, D., Marlot, C. *et al.* Speed of processing in the human visual system. *Nature* **381**, 520–522 (1996).
45. Lamme, V. & Roelfsema, P. The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci* **23**, 571–579 (2000).
46. Hegdé, J. Time course of visual perception: Coarse-to-fine processing and beyond. *Prog Neurobiol* **84**, 405–439 (2008).
47. Stanford, T., Shankar, S., Massoglia, D., Costello, M. & Salinas, E. Perceptual decision making in less than 30 milliseconds. *Nat Neurosci* **13**, 379–385 (2010).
48. Chassy, P. & Gobet, F. A hypothesis about the biological basis of expert intuition. *Rev Gen Psychol* (2011).
49. Guida, A., Gobet, F., Tardieu, H. & Nicolas, S. How chunks, long-term working memory and templates offer a cognitive explanation for neuroimaging data on expertise acquisition: A two-stage framework. *Brain Cognition* **79**, 221–244 (2012).
50. Masunaga, H., Kawashima, R., Horn, J., Sassa, Y. & Sekiguchi, A. Neural substrates of the topology test to measure fluid reasoning: An fMRI study. *Intelligence* **36**, 607–615 (2008).
51. Lee, B. *et al.* White matter neuroplastic changes in long-term trained players of the game of “Baduk” (GO): A voxel-based diffusion-tensor imaging study. *Neuroimage* **52**, 9–19 (2010).
52. Gobet, F. *et al.* Chunking mechanisms in human learning. *Trends Cog Sci* **5**, 236–243 (2001).
53. Simon, H. & Chase, W. Skill in chess: Experiments with chess-playing tasks and computer simulation of skilled performance throw light on some human perceptual and memory processes. *Am Sci* **61**, 394–403 (1973).
54. Charness, N., Reingold, E., Pomplun, M. & Stampe, D. The perceptual aspect of skilled performance in chess: Evidence from eye movements. *Mem Cognition* **29**, 1146–1152 (2001).
55. van Zoest, W. & Hunt, A. Saccadic eye movements and perceptual judgments reveal a shared visual representation that is increasingly accurate over time. *Vision Res* **51**, 111–119 (2011).
56. Muckli, L. What are we missing here? Brain imaging evidence for higher cognitive functions in primary visual cortex V1. *Int J Imag Syst Tech* **20**, 131–139 (2010).
57. Oliva, A. & Torralba, A. Modeling the shape of the scene: A holistic representation of the spatial envelope. *Int J Comput Vision* **42**, 145–175 (2001).
58. De Groot, A. *Thought and choice in chess*, vol. 4 (Mouton De Gruyter, 1978).
59. Waters, A. & Gobet, F. Mental imagery and chunks: Empirical and computational findings. *Mem Cognition* **36**, 505–517 (2008).
60. Nosofsky, R. Attention, similarity, and the identification–categorization relationship. *J Exp Psychol Gen* **115**, 39 (1986).

Author Contributions

M.H. originated the idea and developed the methodology, compiled and analysed the data and wrote the paper. T.B. and A.S. helped interpret the data and write the paper. This work was supported by USAF grant AOARD 104116 and ARC grant DP0881829.

Additional information

Competing financial interests: The authors declare that they have no competing financial interests.

License: This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/3.0/>

How to cite this article: Harré, M., Bossomaier, T. & Snyder, A. The Perceptual Cues that Reshape Expert Reasoning. *Sci. Rep.* **2**, 502; DOI:10.1038/srep00502 (2012).