

This contribution is part of the special series of Inaugural Articles by members of the National Academy of Sciences elected on April 30, 1996.

Geometries of knowledge

(pattern theory/knowledge representation/image inference/digital anatomy)

ULF GRENANDER†

Brown University, Providence, RI 02912

Contributed by Ulf Grenander, December 13, 1996

ABSTRACT It is argued that knowledge representations formalized through pattern theoretic structures are geometric in nature in the following sense. The configurations and resulting patterns appearing in such representations exhibit invariances with respect to the similarity groups and are characterized topologically through their connection types. Starting with a special pattern from microbiology, it is shown how the basic pattern theoretic concepts are introduced in general and what their function is in representing knowledge. Variance/invariance of the patterns is discussed in geometric language. The measures on the configuration spaces are implemented by difference/differential equations which are used as a basis for computer algorithms.

Information and Knowledge. Fig. 1 *Left* is a micrograph of a cardiac cell in a rat. In it several types of structures can be seen, among them some oblong areas, sometimes oval, sometimes curved, the mitochondria of the cell. The mitochondria, whose functions include supplying energy for the cell, can be seen to vary a good deal between each other. Nevertheless they exhibit some stability: their size, while variable from one to another, is of the same order of magnitude, and their texture is characteristic and different from that of the surrounding cytoplasm.

These statements could be made more precise and given quantitative form in terms of lengths, areas, intensities and so on, as is customary in biological/medical research. It is less obvious how to assert that the shape also has some stability, how the concept of shape should be formalized in quantitative form without introducing unrealistic assumptions stipulating a fixed form for all the mitochondria.

This difficulty is not limited to mitochondria, nor to microbiology, since the notion of shape pervades biology in general. In brain research, for example, the anatomies are described through basic units like midbrain, ventricles, hippocampus, . . . , all the way down to the level of cell types as in Brodman's classical brain map. Or, going in the opposite direction, starting with cells, the anatomist builds larger structures which themselves are *combined* into even larger units until the entire organism has been accounted for.

During this building process, in which successive levels are introduced, with units that are measured and located with respect to each other, it is practical to express the descriptive anatomical statements in some coordinate system. It is clear that, *the relations should be invariant with respect to the Euclidean group of transformations.*

A large collection of such micrographs contains a lot of information in raw form about mitochondria, but it is only

when information is organized into ordered systems that it becomes knowledge, say in the form of a textbook, or a theory, a doctrine. To fix ideas we shall limit ourselves to discussing the form of mitochondria, but the reasoning is of much greater generality.

Before entering this discussion it is reasonable to ask why such representations are needed. Say that a biologist has accumulated evidence about some organism, for example in the format of pictures such as the one in Fig. 1. If the number of pictures is small they can be examined "manually" one by one, and the raw data can be reduced to quantitative measurements and summarized into meaningful statements. When the observational techniques become increasingly automated the researcher will be faced by mountains of data and it is reasonable to turn to computational methods of analysis. But to develop computer programs for such a task requires that it be formulated in exact terms: the computer is unforgiving when it comes to nebulous or ambiguous instructions. We are thus forced to formalize our understanding of the research object in well defined logical categories so that they can be correctly translated into computer code: we need, indeed, a *formal representation* of the knowledge.

To represent knowledge by doctrines, perhaps by mathematical models, has a long history going back to the dawn of scientific thinking. Sometimes this has resulted in highly concise theories of great beauty; Newtonian mechanics is the example *par preference*. Others are differently expressed, not using mathematical terminology but nevertheless organizing a mass of information into knowledge representations; the botanical taxonomy of Linnaeus could serve as an example. It is when we are confronted by really complex systems, heterogeneous, highly variable, and with complicated interactions between their parts, that the need arises for formal structures that has led to the creation of *pattern theory*. Biology/medicine offers many examples of this.

Knowledge as Geometry. To be more concrete, let us return to the representation of the shape of mitochondria in micrographs. In Fig. 2*a* we show a polygon whose shape is at least reminiscent of that of a mitochondrion. Let us denote the vertices of an arbitrary polygon, not necessarily closed, see Fig. 2*b*, by $x^i = (x^i_1, x^i_2) \in \mathbf{R}^2$; $i = 1, 2, \dots, n + 1$ and think of the polygon as generated by the n sides

$$g^i = (x^i, \vec{x}^{i+1}).$$

To ensure that the sides are contiguous we introduce *bond values*

$$\beta_{\text{in}}(g^i) = x^i; \beta_{\text{out}}(g^i) = x^{i+1}$$




FIG. 1. (Left) Micrograph of rat cardiac cell. ($\times 20,000$.) (Right) Same section, with mitochondria outlined.

with the obvious bond relation

$$\beta_{\text{out}}(g^i) = \beta_{\text{in}}(g^{i+1}); i = 1, 2, \dots, n - 1.$$

In this illustrative example the choice of bond relation is obvious but that will not always be the case.

Formally we shall write the representation of the n -gon as a configuration

$$c = \text{LINEAR}(g^1, g^2, \dots, g^n),$$

with $\sigma = \text{LINEAR}$ meaning a connector graph, here just a linear graph. These configurations, constrained by the bond relation, form a configuration space $\mathcal{C} = \mathcal{C}(\mathcal{R})$, where the symbol \mathcal{R} stands for the regularity expressed by the bond relation. The constraints are local in the sense that each one involves only a bounded number, in this case two, of generators as $n \rightarrow \infty$. The configuration is interpreted to mean the curve I in the plane corresponding to the polygon and written formally as $R : c \mapsto I; I = Rc$. R is called an identification rule; see below.

We shall also need a subconfiguration space $\mathcal{C}(\mathcal{R}_{\text{closed}})$ consisting of configurations

$$c = \text{CYCLIC}(g^1, g^2, \dots, g^n),$$

where the connector graph *CYCLIC* means that the bond relations are also valid for the bond couple $(\beta_{\text{out}}^n, \beta_{\text{in}}^1)$. We also ask that the polygon not be self-intersecting, which is a global constraint. If that is assumed to be true, Jordan's theorem guarantees that the curve divides the plane \mathbb{R}^2 into two regions, the inside I and its complement I^c . We can then interpret the configuration as "meaning" the image I (see Fig. 2a; we use an identification rule

$$R : c \mapsto I.$$

Note that R is not invertible, since an n -gon encloses a set that can also be enclosed by an m -gon, $m > n$, with some colinear sides. This distinction may at first seem like splitting hairs but is actually fundamental. The representation through a configuration, the *deep structure* to use Chomsky's terminology, is not always uniquely determined by the image, the surface structure. This is like the relation between a function and a formula that represents the function: the formula determines the function but not the opposite.

To handle the permanence of the concept of shape we shall specify a group of transformations that induces invariance. The group will in general be called the *similarity group*, borrowing and extending the meaning of this term from Euclid, and shall act on the space G of generators, in this example

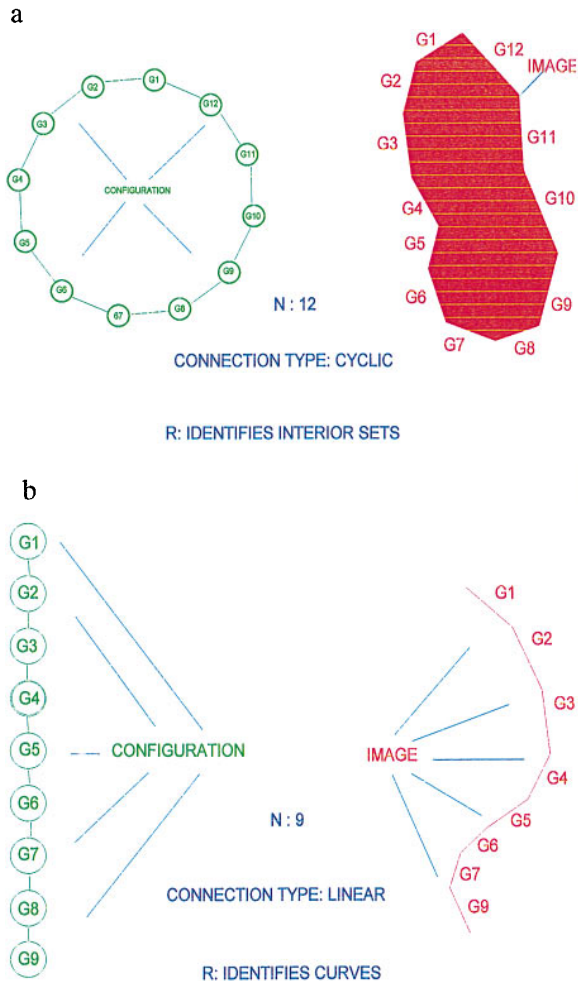


FIG. 2. (a) *CYCLIC*. (b) *LINEAR*.

$$G = \mathbf{R}^2 \times \mathbf{R}^2 = \mathbf{R}^4,$$

since then a generator is determined by its bondvalues, each of dimension two. A good choice of similarity group here has been shown to be the product of the special plane Euclidean group with that of uniform scale change

$$S = SE(2) \times US(2),$$

where a generic group element $s \in S$ acts according to

$$sg = s(\beta_{in}, \beta_{out}) = (s\beta_{in}, s\beta_{out})$$

and extend this definition from G to \mathcal{C} by

$$\begin{aligned} sc &= sLINEAR(g_1, g_2, \dots, g_n) \\ &= LINEAR(sg_1, sg_2, \dots, sg_n). \end{aligned}$$

Note that this transformation acts in the same way on all the generators in the configuration—it is a rigid (up to scale change) transformation. But rigid transformations will not suffice for the representation of flexible shape, since they only describe a single polygon translated, rotated, and uniformly scaled, which is clearly insufficient for describing the variability of mitochondria.

To proceed we introduce local transformations of one or several templates. Starting with a *template configuration*

$$c_{temp} = LINEAR(g_1^0, g_2^0, \dots, g_n^0) \in \mathcal{C}(\mathcal{R}),$$

introduce the *deformed template* by

$$c = LINEAR(s_1g_1^0, s_2g_2^0, \dots, s_n g_n^0)$$

with similarities $s_1, s_2, \dots, s_n \in S$ arbitrary except that they make $c \in \mathcal{C}(\mathcal{R})$. Then the polygon can undergo more flexible changes and we shall see below how to make this a powerful model of shape.

The *image algebra*

$$\mathcal{F} = \{Rc \mid c \in \mathcal{C}(\mathcal{R})\}$$

is indeed an algebraic structure and allows two types of operations. First, the rigid similarities $s : \mathcal{F} \rightarrow \mathcal{F}$ naturally defined through $s : I \mapsto sI$ and, second, combinations of images $I_1 = Rc_1; I_2 = Rc_2; c_1, c_2 \in \mathcal{C}(\mathcal{R})$. Let σ be a connector graph that connects unconnected bonds from c_1 to bonds of c_2 and form the configuration

$$c = \sigma(c_1, c_2).$$

If the new connector is regular, $c \in \mathcal{C}(\mathcal{R})$, so that the last outbond of c_1 agrees with the first inbond of c_2 , we can form the image $I = Rc$ and write formally $I = \sigma(I_1, I_2)$ for the second type of algebraic operation in \mathcal{F} . We also have the sub-image algebra

$$\mathcal{F}_{closed} = \{I \mid I = Rc, c \in \mathcal{C}(\mathcal{R}_{closed})\}$$

representing sets enclosed by polygons as in Fig. 2a. Then we have only one type of algebraic operation, namely $I \rightarrow sI$. In Fig. 3 we show some instances of both types of algebraic operations.

The mapping $R : sc_{temp} \mapsto I_{temp}$ produces a deformed image template from the deformed configuration template. Further, we can expose any image I to rigid transformations and form the set of images

$$P(I) = \{sI \mid s \in S\}.$$

We shall call $P(I)$ a *pattern* and the set

$$\mathcal{P} = \{P(I) \mid I \in \mathcal{F}\}$$

of all patterns can then be seen to be the quotient space \mathcal{F}/S .

The micrograph in Fig. 1a contains several mitochondria and in general we must take into account the fact that the number is not known beforehand. We do this by extending the choice of connectors: we now allow any connector graph σ consisting of a finite number of cyclic graphs, so that we are dealing with the family of graphs

$$\Sigma_{mult} = \cup_{k=0}^{\infty} CYCLIC^k.$$

This leads to the extended configuration space

$$\mathcal{C}(\mathcal{R}_{mult}),$$

which the identification rule R makes into the larger image algebra

$$\mathcal{F}_{mult} = \cup_{k=0}^{\infty} \mathcal{F}_{closed}^k,$$

namely, the union of Cartesian products of the original image algebra, so that k is the number of mitochondria in the image.

The way in which we represent knowledge relative to groups of transformations inducing invariances is reminiscent of Felix Klein's celebrated Erlangen program (1), in which he suggested that geometries should be viewed as systems in which certain relations are invariant with respect to some group of transformations. We can therefore speak of *pattern theory as a geometry of knowledge*, and we can view configurations as geometric objects. Also, the connector graph, which in general can be quite complicated, describes the neighborhood relations of the knowledge representation which gives rise to a *topology of the knowledge representation*. Hence we are faced by problems in geometry, albeit of nontraditional type.

Representing Variability in a Flexible Geometry. The image algebra \mathcal{F}_{closed} has enough flexibility to represent mitochondria shape. Actually, it can represent very complicated shapes quite different from those of mitochondria: it is *too* general and must be restricted to be useful.

To do this we introduce measures on the configuration space and image algebra that will describe *normal variability* of the shapes. One could define such measures directly on \mathcal{C} or \mathcal{F} but it is more natural to do it for the similarities s_1, s_2, \dots, s_n that give rise to the deformed templates, since this reduction brings us closer to the morphogenesis that produces shape. In other words we shall introduce a measure m on S^n with the intent that m will emphasize the shapes that occur most commonly in nature but give little weight to others.

We shall usually choose m to be bounded so that we can normalize it to a probability measure

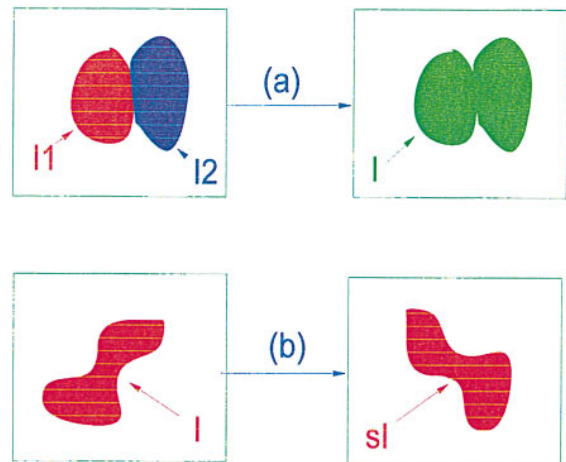


FIG. 3. Algebraic operations. (a) Combination. (b) Similarity.

$$P = \frac{1}{Z} m$$

with the normalization constant $Z = m(S^n)$ known as the partition function to the physicist. The actual construction of m will of course depend upon the particular knowledge domain that is to be represented but should agree with the topology of that domain in the following sense.

For simplicity, say that S is finite-dimensional Euclidean space and that m will be given in terms of a density f on S^n with respect to Lebesgue measure ds^n on S^n . We shall ask that the density, the Radon–Nikodym derivative, have the multiplicative form

$$\frac{m(ds)}{ds^n} = f(s_1, s_2, \dots, s_n) = \prod_{i=1}^n A(s_{i_1}, s_{i_2})$$

with the convention $n + 1 \equiv 1$. The function $A(\cdot, \cdot)$, the *acceptor function*, controls the dependence between the similarity transformations applied to the n generators.

In the special case of mitochondrial shapes, the following specification of the acceptor function has been used successfully. Consider a similarity acting upon a generator $g = (\beta_{in}, \beta_{out}) \in \mathbf{R}^4$ as

$$s : g \mapsto (\beta_{in} + a, \beta_{out} + a) + (rO\beta_{in}, rO\beta_{out})$$

with the translation parameter $a \in \mathbf{R}^2$, the rotation angle $\phi \in [0, 2\pi)$, and the scaling factor $r > 0$, we can read the scaling/rotation part as polar coordinates as

$$rO = r \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix} = \begin{pmatrix} u & v \\ -v & u \end{pmatrix}$$

with $u, v \in \mathbf{R}$. On the four-dimensional similarity group S parametrized by a_1, a_2, u, v we define an acceptor function

$$A(s_i, s_{i+1}) = \exp \left\{ -\frac{1}{2} \left[\frac{Q_{\text{location}}}{\sigma_{\text{location}}^2} + \frac{Q_u}{\sigma_u^2} + \frac{Q_v}{\sigma_v^2} \right] \right\}$$

with the quadratic forms

$$\begin{aligned} Q_{\text{location}} &= (a_1^{i+1} - \rho_{\text{location}} a_1^i)^2 \\ &\quad + (a_2^{i+1} - \rho_{\text{location}} a_2^i)^2 \\ Q_u &= [u^{i+1} - 1 - \rho_u (u^i - 1)]^2 \\ Q_v &= [v^{i+1} - 1 - \rho_v (v^i - 1)]^2. \end{aligned}$$

The rationale behind this choice of $A(\cdot, \cdot)$ is, first, that we want its maximum to be achieved for $a_1 = a_2 = 0$ and $u = 1, v = 0$ implying $r = 1, \phi = 0$, in other words at the identity of the group $S = SE(2) \times SU(2)$. Second, we need to control the variability of the boundary ∂I which we do by means of the variability parameters $\sigma_{\text{location}}, \sigma_u, \sigma_v$ that express how much changes in the location, orientation, and scaling will deviate from the identity. Third, we must also control the smoothness of ∂I through the smoothness parameters $\rho_{\text{location}}, \rho_u, \rho_v$ —by making them small the boundary will be more chaotic and *vice versa*.

This defines a density and measure on $\mathcal{C}(\mathcal{R})$, and we shall restrict it to the sub-configuration space $\mathcal{C}(\mathcal{R}_{\text{closed}})$ by constraining the measure to that manifold, in this example of dimension $2n$ instead of $4n$ for the unrestricted configuration space. Leaving out the technicalities about how to do this, note only that one has to describe the variability of k , the number of mitochondria in the image. This is going to result in a total measure consisting simply of a linear combination

$$m_{\text{multi}} = \sum_{k=0}^{\infty} c_k \underbrace{m \times m \times \dots \times m}_{k \text{ times}}$$

unless we want to exclude the possibility of two mitochondria overlapping each other, which will not be discussed here.

But how does this actually work, how well does the knowledge representation mirror the shape properties of real mitochondria? This question can be answered in two ways, one of which is *pattern synthesis*. If we use Monte Carlo simulation of the representation with appropriate values plugged in for the variability parameters, how do the images look? Fig. 4 shows an empirical template in green and some synthesized shapes in red; they seem to be at least qualitatively like the ones in Fig. 1. But this is not the deciding test for whether to accept or reject the representation. Instead we will have to study how it behaves when applied to real micrographs, *pattern analysis*, and this will be discussed below in *Differential Equations for Patterns*. But first I shall put the previous in a more general perspective.

Geometry of Patterns in Complex Systems. The previous discussion was all in a very special case, but the treatment of patterns in a general system can be based on the same concepts and reasoning. The abstract formulation, abbreviated by omitting some topics that are important but cannot be treated here, reads as follows; more precise statements and proofs can be found in Grenander (ref. 2, chapters 1 and 2 and part IV, and ref. 3).

Starting from a generator space G of primitive objects equipped with a number, the *arity* $\omega(g)$, of bond values

$$\beta_j(g); j = 1, 2, \dots, \omega(g)$$

with values in some bond value space B , and a *bond relation*

$$\rho : B \times B \rightarrow \{\text{TRUE}, \text{FALSE}\}$$

we form configurations

$$c = \sigma(g_1, g_2, \dots, g_n)$$

by putting the generator g_i at the i th site of the connector graph σ from some family, the *connection type*, Σ , of connector graphs. A group S of transformations $s : G \rightarrow G$ acts on G . In most cases that have been analyzed S is a finite-dimensional Lie group. The configuration c is said to be \mathcal{R} -regular with the regularity

$$\mathcal{R} = \langle G, S, \rho, \Sigma \rangle$$

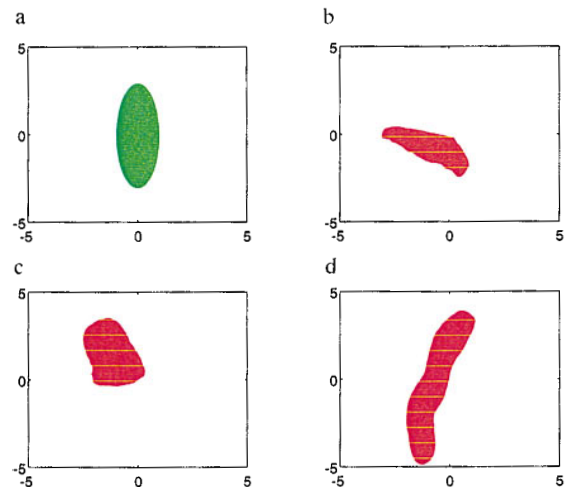


FIG. 4. (a) Template. (b–d) Synthesized shapes.

if all pairs of bonds emanating from two sites (i_1, i_2) connected by σ satisfy the bond relation

$$\rho(\beta_{j_1}(g_{i_1}), \beta_{j_2}(g_{i_2})) = TRUE$$

and $\sigma \in \Sigma$. See Fig. 5. The regularity \mathcal{R} leads to a configuration space $\mathcal{C}(\mathcal{R})$ and, with some equivalence relation R , the identification rule, we get the image algebra as the quotient

$$\mathcal{F} = \mathcal{C}(\mathcal{R})/R = \mathcal{C}(\langle G, S, \rho, \Sigma \rangle)/R$$

and the pattern family

$$\mathcal{P} = \mathcal{F}/S.$$

The regularity \mathcal{R} determines the geometry of the knowledge represented by the *regular structures* $\mathcal{C}(\mathcal{R})$ and together with the identification rule R it determines also the algebraic structure of \mathcal{F} .

Returning to the configuration space, we introduce a measure m as a Radon-Nikodym derivative with respect to a σ -finite measure μ^n on S^n through the *structure formula*

$$\frac{m(ds^n)}{\mu(ds^n)} = \sum_{\sigma \in \Sigma} c_\sigma \times \prod_{((i_1, j_1), (i_2, j_2)) \in \sigma} A(\beta_{j_1}(g_{i_1}), \beta_{j_2}(g_{i_2}))$$

with some acceptor function

$$A : B \times B \rightarrow \mathbf{R}_+$$

and the product taken over all pairs $(i_1, j_1) \leftrightarrow (i_2, j_2)$ connected by the graph σ . When S is a Lie group it is natural to choose the measure μ^n as an invariant measure on the group. The measure m induces measures on the image algebra \mathcal{F} and patterns in \mathcal{P} just as in the mitochondria example. If the measure m is bounded it can be made into a probability measure by normalizing with the partition function $Z = m(\mathcal{C}(\mathcal{R}))$. These measures determine the metric structure of the knowledge representations.

Differential Equations for Patterns. Once we have acquired some understanding of the geometry of a pattern class it can be exploited for practical use. It then depends upon what sensor was used to acquire the picture. Medical sensor technology has advanced rapidly in recent years and resulted in a multitude of modalities—for example, magnetic resonance, ultrasound, positron emission tomography, and visible light microscopy to mention but a few. The micrographs in the first section of this paper were obtained by an electron microscope, and we shall discuss the analysis of them, pointing out that the

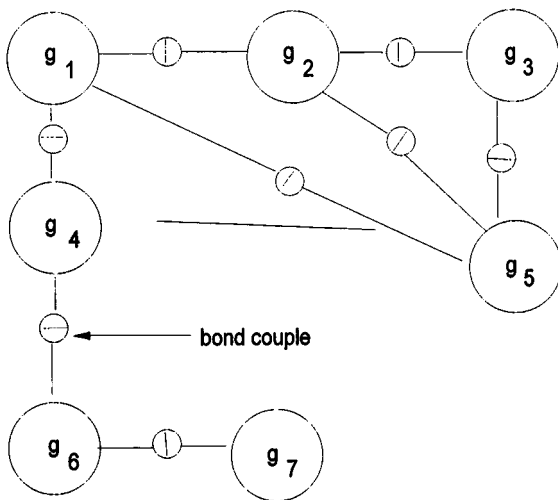


FIG. 5. Configuration diagram.

mathematical method is essentially the same for all modalities except, of course, that the physics can differ radically from one sensor to another.

Call the observed image deformed by biological variability as well as by sensor imperfections $I^\mathcal{D}$. In most cases the deformed image can be represented by an array so that we can write

$$I^\mathcal{D} = \{I_y^D; y \in Y\},$$

where Y stands for the array and its entries of $I^\mathcal{D}$ are real numbers. One can make this concrete by thinking of Y as parametrizing the “photographic plate” taking this term in a wide sense. The biological variability is described by a prior probability density $\pi(c)$ as discussed in the previous section, and we now have to describe the sensor physics by another probability density $L(y|c)$ conditioned by the statement that $c = s c_{\text{temp}}$ is the true configuration. The form of $L(\cdot|c)$ will vary from sensor to sensor, and I only offer a simple example when Y is a square matrix and I is a subset of the square $(0, l)^2$

$$Y = ((i, j); i, j = 1, 2, \dots, l)$$

$$L(I^\mathcal{D}|c) \propto \exp -\frac{1}{2\sigma_{\text{obs}}^2} \left[\sum_{\text{inside}} (I_{ij}^\mathcal{D} - m_{\text{in}})^2 + \sum_{\text{outside}} (I_{ij}^\mathcal{D} - m_{\text{out}})^2 \right]$$

with the sums extended over the subsets of the $l \times l$ lattice

$$\begin{aligned} \text{inside} &= \{(i, j) | (i, j) \in sI_{\text{temp}}\} \\ &= \{(i, j) | (i, j) \in sRc_{\text{temp}}\} \\ \text{outside} &= \{(i, j) | (i, j) \notin sI_{\text{temp}}\} \\ &= \{(i, j) | (i, j) \notin sRc_{\text{temp}}\}. \end{aligned}$$

The typical intensity inside an object is m_{in} and outside m_{out} . The resolution of the sensor is measured by the value of l and its accuracy by σ_{obs} .

Introduce the functions of s^n , called energies because of their role similar to that of real energies in physics,

$$\begin{aligned} E_{\text{prior}}(s^n) &= -\log \pi(c) = -\log \pi(s^n c_{\text{temp}}) \\ E_L(s^n) &= -\log L(I^\mathcal{D} | s^n c_{\text{temp}}) \\ E(s^n) &= E_{\text{prior}}(s^n) + E_L(s^n). \end{aligned}$$

In this approach *all the knowledge about the unknown true configuration c is contained in the prior knowledge represented by $\pi(\cdot)$ and the empirical knowledge represented by $I^\mathcal{D}$ and can be summarized by the posterior density*

$$p(s^n | I^\mathcal{D}) \propto \pi(s^n c) L(I^\mathcal{D} | s^n c).$$

This is just Bayes’ theorem when the deformed image $I^\mathcal{D}$ is treated as observed and fixed. The problem of how to make strong statements about the pattern represented by the unknown s^n , c , or I is therefore reduced to handling $p(s^n | I^\mathcal{D})$.

To do this we shall employ the differential equation

$$ds^n(t) = -\nabla E(s^n) dt + W(dt)$$

involving the gradient $\nabla E(s^n)$ and the Wiener process $W(\cdot)$ on the Lie group S . The variable t denotes algorithmic time.

This differential equation can be motivated in two ways. The first, simple but superficial, is to point out that the first term on the right makes the trajectory through S^n move against the gradient of the energy $E(s^n)$ so that it tends toward a local energy minimum. Many systems in the physical sciences are

controlled by such a behavior, so it is not unreasonable to postulate this sort of strategy for interpreting an observed image.

A deeper motivation is to consider the probability measure $P(t)$ of $s^n(t)$ and to note the well known fact that

$$P(t) \xrightarrow{\text{weakly}} P$$

as $t \rightarrow \infty$, where P is the measure with density $p(s^n|I^{\otimes n})$. That means that if we solve the differential equation from $t = 0$ up to a large value t_{final} , then the resulting value

$$s^n(t_{\text{final}}) \underset{\approx}{\text{in measure } P},$$

so that $s^n(t_{\text{final}})$ is a likely *explanation* of $I^{\otimes n}$. Doing this a number of times will give us a sense of what accuracy we can attach to the explanation. Many variations of this statement can be made, but the above one will have to suffice here. We shall bypass a technical difficulty, ergodicity, here and only refer to Grenander and Miller (4).

Of course the variability of k , the number of mitochondria in a micrograph, must also be taken into account. Without entering into a discussion of how this has been done we just mention that the differential equation will be extended to a difference-differential equation for which the solution is allowed to jump at discrete time points corresponding to the decision $k \rightarrow k \pm 1$ and follow the differential equation between the jumps. Fig. 6 shows schematically how the solution develops through configuration space.

In the mitochondria study the output of the differential-difference equation can look like Fig. 1 *Right*, which shows the boundaries of the recognized mitochondria.

Let us now turn to a more challenging task, to represent variability in images obtained by magnetic resonance, positron emission tomography, or single photon emission computed tomography. Here too the formulations should start in the Euclidean continuum, not on the lattice of the observed picture as has often been done in traditional pattern recognition, which is one of the reasons for its difficulties. We are now in \mathbf{R}^3 . Configurations are then built from generators consisting of the components of the brain, hierarchically organized, and with bondvalues made up from surface patches. Further, the generators may carry parameters—for example, indicating tissue type or the name of the component. The connectors, the σ s, that represent the topology of anatomical atlases, say for normal anatomies, express contiguity of surface patches.

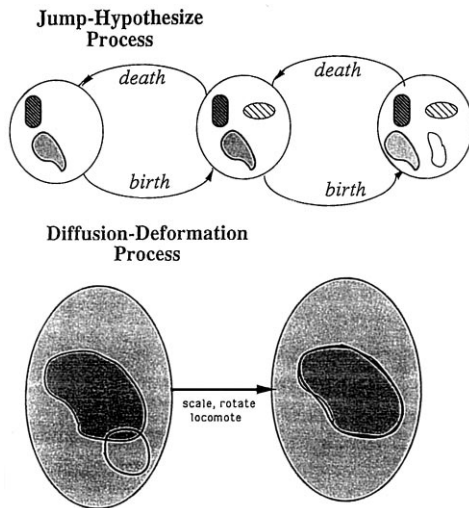


FIG. 6. Solution development.

For (normal) anatomical patterns the similarity transformations will be made up of a *cascade of groups*. On each group in the cascade a measure is introduced to represent that part of the variability that is due to the transformations in the group. This leads to a density $p(c|I^{\otimes n})$ by an argument that is in principle the same as the one above, although more complicated, and also to differential equations over the respective groups.

Given a template brain represented by a configuration

$$c_{\text{temp}} = \sigma(g_1, g_2, \dots, g_n) \in \mathcal{C}(\mathcal{R})$$

we acquire for a patient an observed image $I^{\otimes n}$ and then solve the differential equations to find a reasonable diffeomorphism $d^* : I_{\text{temp}} \rightarrow \mathcal{F}$ that will serve as an explanation of the observation. We then let d^* operate on the generators in the image template

$$\begin{aligned} d^*I_{\text{temp}} &= d^*Rc_{\text{temp}} \\ &= R^*\sigma(g_1^*, g_2^*, \dots, g_n^*); g_i^* = d^*g_i. \end{aligned}$$

This makes it possible to *achieve automatic recognition of the components of the patient brain as Rd^*g_i* . Fig. 7 indicates the nature of such automatically obtained mappings, or brain warps. Once the components have been identified it is a simple matter to let the computer calculate their volumes, areas, and curvatures, as well as other geometric characteristics that may be of help to correlate with the symptoms and signs observed for the patient.

Of course this procedure is justified only for a patient population with brains topologically equivalent to I_{temp} . With conscious abuse of medical terminology we name a brain normal only if it is topologically equivalent to the template in the given sub-population; this we consider a necessary (but not sufficient) condition for normality. How to deal with images of abnormal brains requires an extension of the methodology described above, but we cannot go into this challenging topic here.

Geometric Understanding of Patterns. We have argued that pattern theoretic representations of knowledge are inherently geometric in nature. Their geometry is described metrically by transformations from Lie groups, the similarities S . They are described topologically by graph structures Σ and jump transformations acting upon these connection types. With the aid of such *variability transformations questions of inference from observed images reduce to the geometric understanding of the pattern*. The two examples discussed above happened to be pictorial; this was completely incidental, however, and the statement is not limited to understanding of observed pictures.

Such understanding is possible for patterns of general nature by solving the differential equation that governs the system and is derived from a pattern theoretic representation. The aim of these knowledge representations is fundamentally different from that in artificial intelligence, AI, in which the same term occurs. In AI the overall objective is to imitate human intel-

Deformable Brain Mapping




FIG. 7. Deformable brain mapping.

ligence in general by computer algorithms. The goal of pattern theory, on the other hand, is at the same time more limited and more ambitious. It is limited in that it does not pretend to be able to imitate general intelligence, only narrow segments of human mental activities such as recognition of mitochondria or anatomical components. It is ambitious in that it tries to automate such activities so that the task can be carried out faster and perhaps better than by humans.

But this approach is also limited to really complex systems, where it is important to realize that “complex” does not just mean big but has a more restricted meaning. A system with lots of components, such as an ideal, monatomic gas, is certainly complicated, and its study has raised challenging mathematical questions. Nevertheless the rules that are assumed to govern its behavior are straightforward and no pattern formalism is needed—it would help little if at all. In contrast, systems in biology/medicine often exhibit behavior characterized by a

high degree of heterogeneity, complicated interactions, and awesome variability. Then a formal representation helps to separate the essential from the incidental and sets the mathematical stage for a rational approach to automated image understanding in the limited sense above. The last decade has witnessed the successful application of this methodology to microbiology, digital anatomy, language theory, and automatic target recognition, and some attempts are being made at present to use these ideas in nonnumerical situations.

1. Klein, F. (1893) *Math. Ann.* **43**, 63.
2. Grenander, U. (1993) *General Pattern Theory* (Oxford Univ. Press, Oxford).
3. Miller, M. I., Christensen, G. E., Amit, Y. & Grenander, U. (1993) *Proc. Nat. Acad. Sci. USA* **90**, 11944–11948.
4. Grenander, U. & Miller, M. I. (1994) *J. R. Stat. Soc. Ser. B* **56**, 549–603.