A Holarchical Approach to Robot Vision

David Vernon

Department of Computer Science,
Trinity College, Dublin, Ireland.

Abstract

This paper describes initial progess toward the development of a holarchical robot vision system. The holarchical concept is an organisational philosophy which is borrowed from system sciences and is concerned with the organisation of interacting complex systems. The major levels of this holarchy correspond to the peripheral, attentional, and cognitive visual processes. The approach is exemplified by describing the development of a contour-following boundary-based object recognition vision system. The extension of the holarchy to incorporate further visual cues is considered.

Introduction

The success of man's endeavours to comprehend and utilise his environment, exemplified by three centuries of significant technological developments, lies deep-rooted in reductionist analysis, decomposing systems into their constituent processes or functions, and thus enabling the problem to be comprehended and theories for its explanation or solution to be advanced. Contemporary research in information technologies appears to be based almost entirely on the premises of reductionism. The belief implicit in reductionist philosophy is that the system (whatever it is) may be understood by a complete and rigorous analysis of these parts and relations. Unfortunately, it appears that many complex (and, in particular, biological) systems do not lend themselves to this approach and several eminent scientists from diverse fields have attested to this inadequacy in reductionism[1]. The dilemma may be subtly summarised by the observation (attributed to Jan Smuts) that the whole (i.e. the system) is greater than the sum of the parts. It is important to note, also, that no trite addition of a 'glueing factor' in terms of inter-component relations or ad hoc fudge factor will balance this inequality. This observation regarding the importance of the wholeness, as an essential characteristic of the system, is the basic tenet underlying the philosophy of holism. Holistic arguments are worthy of our attention, if only because they explicitly acknowledge the fundamental importance of the quality of 'wholeness'. Unfortunately, there has been a tendency among holists to use this wholeness or gestaltness in a somewhat axiomatic manner, i.e., as something which does not require further explanation.

The central theme of this paper in not so much to report, a posteriori, the results and performance of a vision system, but to propose and argue the necessity of an alternative approach, i.e., that the valid aspects of both philosophies of reductionism and holism must be combined if one is to address and understand complex systems and, in particular, visual perception. A central and significant problem is posed by the necessity to utilise holistic phenomena in a concrete manner; it is proposed here that the concepts of holons and holarchies introduced by Koestler[2] should be adapted and used as the link between reductionism and holism applied to the understanding of vision.

The term 'holon', originally coined by Koestler from the Greek 'holos' (for whole) and the suffix 'on' (implying particle or part as in proton), is intended to convey an entity which is a sub-whole, i.e. "a stable integrated structure equipped with self-regulatory devices, and enjoying a considerable degree of autonomy or self-government"[2]. More specifically, the term may be applied to any structural or functional sub-system in a biological, social, or cognitive hierarchy which manifests rule-governed behaviour and/or gestalt constancy. It is frequently asserted that hierarchic organisation is an essential and distinguishing characteristic of sociological and biological phenomena.[2,3,4] Weiss emphasises that such organisation is a real phenomenon, presented to us by the biological object and is not the fiction of a speculative mind.[5] Koestler argues that the hierarchy is a natural manifestation of a holistic organisational structure and, in particular, it facilitates "the emergence on each higher level of new 'organising relations' between (sub) wholes of greater complexity, whose properties cannot be reduced to, not predicted from, the lower level"[2]. Such hierarchies are constructed at each level of entities which are, to an extent, self-sufficient (and hence 'wholes') and simultaneously constituent parts of the overall hierarchical structure. The holon fits this description most aptly and such a hierarchy of holons is referred to as a holarchy. Thus, the holarchy is a hierarchy of Janus-faced entities (holons), the upward face representing the dependent aspect which is part of the hierarchy above it and the downward face representing the autonomous aspect tending to self-sufficiency.

A caveat is in order, however. While hierarchic organisation is usefully exploited in many areas, it would be inappropriate to suggest that holarchies are merely, and nothing but, hierarchies and that the concept is, thus, trivial. Holarchies comprise entities with distinct and special properties, which facilitate the concrete emergence of holistic characteristics, especially since holons are quasi-autonomous entities and function in parallel. The complete system represented by this holarchy cannot be described by a purely reductionist analysis of each separate holon in isolation. Only by functioning in 'coalescence' with the others does the holon reveal its true character (see also ref. 6). Further, the concurrent operation of the individual holons implies a dynamic system (facilitated by hierarchic structure), the global operation of which is inherently inaccessible to a strictly functional decomposition. The following text discusses in more detail the properties and characteristics of these holons and holarchies.

Holons and Holarchies

Every level in a holarchy of any type is governed by a set of fixed invariant rules which account for coherence, stability, and the specific structure and function of its constituent holons, referred to as its 'code' or 'canon'. It is important to note, however, that although these canons are invariant, they allow for flexible strategies. Thus, in simple terms, the canons correspond to the rules of the game, while leaving the actual strategy adopted during play to be determined by other factors, typically the holon's local environment (which may well include interaction with other holons). Koestler asserts the existence and fundamental importance of two polar tendencies inherent in holons, which must be made explicit in the holon's canons. These are the holon's integrative (or self-transcending) tendencies to function as part of the larger whole and its self-assertive tendency to preserve its individual autonomy. Thus: "the self-assertive tendency is the dynamic expression of the holon's wholeness while its integrative tendency is the dynamic expression of its partness"[2]. Several general characteristics of these holarchies are detailed. The first of these, the characteristic of dissectibility, is a natural expression of the structure implicit in the holarchy. Holarchies are dissectible into their constituent branches, the nodes of which are the holons, and the hierarchic connecting lines represent the channels of communication and control. One should interpret the word 'dissectible' cautiously for to interpret it as being divisible would be to do so in a purely reductionist sense. This would clearly be a tacit refutation of the holarchic philosophy and it is preferred to attach a meaning closer to 'distinguishibility'[6]. Thus, the distinctiveness of the holons is explicit but in no way can they be separated (in the reductionist sense) as they are, by definition, mutually dependent, on the grand holarchic scale, at least.

A second property of holarchies, and a property possessed by all hierarchies, concerns the vertical modes of interaction between levels in the holarchy. Specifically, vertical communication is only allowed on a level-to-level basis, through 'regulation channels'. There is no facility for short-circuiting the communication process. This strict vertical connection and communication is referred to as 'arborization', a term which derives from the inherent branching structure of the holarchy. There is, however, a complementary principle to arborization, that of 'reticulation'. Holarchies do not exist in an isolated void: frequently there are other holarchies, often with a closely related purpose, which may exist or be operating simultaneously. That there may be communication between constituent holons of each holarchy is explicitly acknowledged; this process is referred to as reticulation or net-formation. Reticulation may exist on one or on several levels simultaneously. Koestler suggests that the perceptual and motor systems of man form two distinct holarchies exhibiting reticulation at several levels.

Holons and holarchies do possess other characteristics, such as a regenerative ability to adapt and evolve by ordered restructuring, but these are not of direct relevance to this discussion and will be neglected.

Perceptual Organisation

As mentioned above, no formal methodology exists to aid the identification of the constituent holons of the holarchy or the terms of the system. With specific regard to their identification in the context of visual organisation, one is forced to hypothesise a feasible structure based on existing research in psychology and machine perception. Martin and Aggarwal in a review article on dynamic scene analysis indicate that much research in computer vision has followed the human perceptual divisions comprising cognitive, attentive, and peripheral processes.[7] Attentive processes operate on stimuli which impinge on the fovea, peripheral processes operate on the periphery of the visual field, "watching" for interesting features such as motion, texture, and colour, and direct the attentive processes to it: the the attentive processes must be able to track to movement and attend to the detail of the objects in motion. Cognitive processes decide which area of interest should next be attended to and relates peripheral and attentive processes to the person, his knowledge, and expectations. Amore recent dynamic scene analysis system VILI[8] is also organised according to this three-level hierarchy and incorporates knowledge sources at each of these levels. The knowledge sources at each level contain information pertinent to analysis at that level alone, e.g. the knowledge source at the cognitive level would contain image independent general knowledge whereas the peripheral level would be able to access image-specific information but not general information. Pau identifies a similar hierarchical architecture comprising cognitive, attentive, and feature extraction processes.[9] The importance of attentive processes to segmentation and analysis problems[10,11] provide further motivation to incorporate these divisions in an organisation or architecture which recognises explicitly their relevance and relationship to other visual processes.

The research described in this paper is based on the assumption that such a perceptual organisation is, indeed, valid. Note that the original concept of peripheral processes has been extended to include, in a general way, the extraction of all relevant visual cues, at several resolutions, from the visual image. Further, the system has been conceived as a general visual perception system and made explicit in the context of a contour-based dynamic boundary-following system to facilitate initial implementation and to gain a greater understanding of the approach. A prior attempt at utilising this perceptual organisation in a true hierarchic manner as a boundary-following system[12] functioned very satisfactorily with simple scenes but attempts to extend it in a coherent manner to incorporate other visual cues were abortive. This, then, provided significant motivation for the holarchic approach propounded here. This (initial) visual holarchy, then, comprises three levels. At the top level, or apex, there exists just one holon, the holon of visual perception. There are three holons at the second level: these are the cognitive, attentional, and peripheral holons. It should be noted that this tiered structure is slightly misleading; the holon of visual perception encompasses and comprises the three holons at the level under it, thus a paradigm of nesting rather than tiering would be more appropriate. Each of these three holons comprise their own sub-holarchy. The cognitive holon comprises, at level three, two holons of recognition and of inference and modelling. The attentional and peripheral holons comprise a holon for each distinct visual cue, e.g., edge, motion, stereo disparity, and shading. Note that in the contour-following system there is just one cue, that of edge data. Each of these holons must be fully identified by the specification of the canons that govern their behaviour; the remainder of this paper details the canons of function, structure, integrative tendency, and self-assertive tendency, together with a specification of the information representations associated with both integrative and self-assertive tendencies. Due to space constraints, only the visual perception, cognitive, attentional, and peripheral holons are dealt with in this paper; the canons and information representations of the remaining holons may be obtained by contacting the author.

A specification of canons and information representations for a holarchic robot vision system

Holon of visual perception: canon of function

General statment. This holon shall infer or hypothesise the three-dimensional (3-D) structure of a perceived visual scene from binocular stereo images representing the reflectance function (incident light intensity), assigning some episematic label based on prior perception which has been validated by positive feedback.

Contour-specific statment. This holon shall infer or hypothesise the 2-D structure (shape) of objects within a perceived visual scene from a single image representing the scene's reflectance function and shall assign an episematic label based on prior perception which has been validated by positive feedback. The 2-D structure/shape shall be based exclusively on the object's boundary.

Holon of visual perception: canon of structure

General statment. This holon shall be considered to be the apex of the holarchy of visual perception and, at this stage of development, no interaction is assumed with other perceptual faculties or holarchies, e.g. aural and tactile senses. Further, it will be assumed that there is no interaction with the motor system. (Note: this assumption is restrictive and may have to be relaxed to accommodate saccadic eye-movements and head movement associated with attention and visual motion.) The holon is organised as three sub-holons of attentional, cognitive, and peripheral processes.

Contour-specific statment. As with the above general statment, the holon is, in effect, an isolated (self-assertive) and self-sufficient module, situated at the apex of the visual holarchy. Thus, all inferences made regarding the external environment are made solely on the basis of visual information. Again, the holon is organised as three sub-holons of attentional, cognitive, and peripheral processes, which are detailed in the canon of the self-assertive tendency.

Holon of visual perception: canon of integrative tendency

General and contour-specific statements. Note: since one is assuming that this holon is the apex of a distinct and independent holarchy, it follows that the integrative tendency is minimal. However, the following, at least, must be asserted. The holon of visual perception shall in all its derived conclusions, representations, and operation strive to be consistent with current context and expected situations, based on prior experience (learned perception) or a priori knowledge (embedded information). This shall be accomplished such that the least amount of reconciliation between discrepancies is required. (It is not possible to formalise this completely as to do so would require the (holarchic) identification of cognition and this is not yet within the present frame of reference.)

Holon of visual perception: canon of self-assertive tendency

General statement. The visual perception apex holon shall function so that all its derived conclusions (i.e. inferences and hypotheses), representations, and operations are mutually, and internally, consistent. Thus, the 3-D representations must be derived via the simplest path on the basis of available visual cues. The operation of the holon is based on the mutually co-operative function and interaction of three sub-holons, corresponding to cognitive, attentional, and peripheral processes. (Canons of the self-assertive tendency must address the coodination of the activity of these three holons, thus:)

$A|^{VP}|C$ Cue information generated by the attentional holon shall be made available to the cognitive holon.

$C|^{VP}|A$ Quantitative assessment of the consistency of information generated by the attentional holon with prior experience (or a priori knowledge), represented by models, will be made available by the cognitive holon to the attentional holon through the visual perception holon.

$P|^{VP}|A$
$P|^{VP}|C$ Representations of visual cues generated by the peripheral holon shall be available to both the attentional and the cognitive holons.

$A|^{VP}|P$ The spatial extent (which is a function of operating resolution) and the cue emphasis of peripheral processes shall be directed by the attentional holon.

$C|^{VP}|P$ Peripheral processes do not take cognisance of the operation of the cognitive processes (or, rather, no meaningful interpretation of such interaction can be envisaged at this time).

$C|^{VP}$ The cognitive holon shall provide the visual perception holon with a fully identified canonical object-centred model of the environment, together with an episematic label.

Contour-specific statment. All the above canons apply. However, the following restrictions are required.

$A|^{VP}|P$ Spatial extent is fixed and only one resolution shall be used. Furthermore, cue emphasis is redundant since the peripheral holon utilises only one cue (edges or intensity discontinuities).

$C|^{VP}$ The canonical model reduces to a 2-D shape descriptor, fully identified (position and orientation) and including an episematic label.

$C|^{VP}|A$ Contours must be consistent with expected shape; quantitative assessment will be provided regarding such consistency.

Information representations associated with integrative tendency

General case. A canonical 3-D object-centred model of the perceived environment, together with episematic labels (episemons). Note: such a representation might be based on generalised cylinders.

Contour-specific case. A boundary-based shape descriptor, in particular a NCD (Normal Contour Distance) descriptor, completely identified, i.e., including position, orientation, and episemon. Homogeneous transformations may be used for position and orientation specifications. The camera model will facilitate object-centred description.

Information representations associated with self-assertive tendency

General.

A|VP|C — Higher-level image-centred scene representations, typically of the form of the primal sketch, i.e. coherent semi-local 2½-D representations based on several visual cues.

C|VP|A — A consistency measure between the above A|VP|C representations and the expected models.
A set of spatial deviations which may offer more successful model consistency.

P|VP|C
P|VP|A — Image-based representations of all low-level cues derived by the peripheral level.

A|VP|P — A measure of the spatial extent (assuming a regular geometric support, e.g., circle or square), centred on the current attentional focus, and specified for each operational resolution.
An indication of cue emphasis (which may be interpreted as priority by cue and/or priority by resolution).

C|VP|P — No meaningful representation has yet been identified.

C|VP — Object-centred canonical model of the environment, fully identified, and including episematic labels.

Contour-specific.

A|VP|C — Image-centred boundary-based shape representation, based exclusively on the edge cue.

C|VP|A — A consistency measure between the above A|VP|C representations and the expected models.
A set of spatial deviations which might offer more successful model consistency.

P|VP|A
P|VP|C — Image-based representations of edge cue derived by the peripheral holon.

A|VP|P — The spatial extent is trivial in this case and is confined to an elementary 3x3 neighbourhood, centred on the current attentional focus.
Cue emphasis is also trivial: there is only one cue and only one resolution.

C|VP|P — No meaningful representation has yet been identified.

C|VP — NCD 2-D shape descriptor, fully identified, and including an episematic label.

Holon of peripheral processes: canon of function

General statement. This holon extracts all of the several distinct visual cues, e.g. edges or intensity discontinuities, stereopsis, motion. (The best of available theories of low-level vision are used; in particular, those in the vein of Marr, Poggio, Ullman, and Hildreth). Hence, the processes should operate in an efficient manner, perhaps co-operatively, using as much information as possible from mutually relevant cues, i.e., low-level modules.

Contour-specific statement. This holon extracts just one cue: edges or intensity discontinuities.

Holon of peripheral processes: canon of structure

General statement. Processes are polar transformations from simpler, more implicit, representations to richer, more explicit, representations. Processes may co-operate in parallel, non-iteratively, and each process is organised as a sub-holon.

Contour-specific statment. No co-operation is required as it is assumed that there is just one cue. This process is, however, still organised as a sub-holon (to facilitate future extensions of the system to incorporate other cues).

Holon of peripheral processes: canon of integrative tendency

General statement.

$A|^{VP}|P$ The domain of activity of the peripheral holon - call it the spatial extent (which is a function of the operational resolution) - and the cue emphasis shall be directed by the attentional holon via the holon of visual perception.

$C|^{VP}|P$ The peripheral holon is not affected by the cognitive holon (or, rather, no meaning can yet be attached to such interaction).

$P|^{VP}|C$ Cue representations are available to both the attentional and cognitive holons.

$P|^{VP}|A$

Contour-specific statement.

$A|^{VP}|P$ The spatial extent is fixed and cue emphasis is trivial; there is only one cue.

Holon of peripheral processes: canon of self-assertive tendency

General and contour-specific statement. Processes may operate in parallel on all spatial extents and at all resolutions, co-operating on a non-iterative basis, i.e., as sub-holons. Note: in general, high resolution is effective at the centre of the current attention only; it remains the responsibility of the attentional holon, however. Since in-depth knowledge of one cue only is available at present, self-assertive canons are restricted to the associated holon. Thus:

$^{P}|E$ The peripheral holon simply passes on cue emphasis and spatial extent restrictions to the edge holon.

$E|^{P}$ The edge holon makes edge representations available to the peripheral holon.

Information representations associated with integrative tendency

General and contour-specific.

$P|^{VP}|C$, $P|^{VP}|A$, $A|^{VP}|P$, and $C|^{VP}|P$ representations are as detailed in the information representations associated with the self-assertive tendency of the holon of visual perception.

Information representations associated with self-assertive tendency

General.

$E|^{P}$ Cue representations - image-based representation of edges or intensity discontinuities.

$^{P}|E$ A measure of spatial extent and cue emphasis as in $A|VP|P$.

Contour-specific.

$E|^{P}$ Cue representations - image-based representation of edges or intensity discontinuities.

$^{P}|E$ Spatial extent as detailed in the contour-specific information representation $A|VP|P$ associated with the integrative tendency of the peripheral holon.

Attentional holon: canon of function

General statement. The attentional holon shall direct the (spatial) activity of the peripheral holon on the basis of the several distinct visual cues made explicit by the peripheral holon. The attentional holon shall build higher-level image-centred scene representations (typically of the form of the primal and $2\frac{1}{2}$-D sketches, i.e., coherent $2\frac{1}{2}$-D semi-local representations based on several intrinsic images / visual cues.

Contour-specific statement. The attentional holon shall direct the activity of the peripheral holon on the basis of edge cue only. The attentional holon shall build an image-centred boundary-based representation, in particular, an NCD.

Attentional holon: canon of structure

General and contour-specific statement. The holon comprises several attentional processes corresponding to the several visual cues: these are organised as sub-holons. In the contour-specific case, the holon comprises just one sub-attentional process - edge cue attention.

Attentional holon: canon of integrative tendency

General statement.

VP A| |C	The attentional holon shall provide information to the cognitive holon, via the holon of visual perception, regarding shape and cues in the form of higher-level image or scene representations.
VP P| |A	It shall use the cues made explicit by the peripheral holon to construct these representations.
VP C| |A	Derived representations shall be consistent with expectation as indicated by cognitive holon assessment.
VP A| |P	Spatial extent and cue emphasis of peripheral processes shall be directed by this holon.

Contour-specific statement.

VP A| |C	The attentional holon shall provide information to the cognitive holon, via the holon of visual perception, regarding shape using a boundary-based (partial) shape descriptor.
VP P| |A	It shall use the edge data made explicit by the peripheral holon to construct these representations.
VP C| |A	Derived shape representations (contours) shall be consistent with expected shape: quantitative assessment will be provided regarding such consistency by the cognitive holon.
VP A| |P	Cue emphasis is redundant - one cue only is used. Spatial extent is trivial - it is confined to the minimum local detection region, e.g. a 3x3 neighbourhood.

Attentional holon: canon of self-assertive tendency

General statement. As these canons must be based on evolving theories of integration of low-level modules, discussion is restricted here to attentional processes in the contour-specific case.

Contour-specific statement.
- Contours comprise continuous locally maximal gradient segments (or connected points of zero-crossing of an image formed by the convolution of the intensity image with a Laplacian of Gaussian mask).
- Contours do not change direction discontinuously (without reason). If this occurs, one must look for alternative contours and/or evidence that a discontinuous contour is correct using other cues (if available).
- Evidence must be sought on a semi-local basis.
- Different (lower) resolution representations will be used to guide activity on higher resolutions, especially for semi-local activity.

Information representations associated with integrative tendency

General.

$A|^{VP}|C$, $C|^{VP}|A$, $P|^{VP}|A$, and $A|^{VP}|P$ representations are as detailed in the general information representations associated with the self-assertive tendency of the holon of visual perception.

Contour-specific.

$A|^{VP}|C$ — Image-centred boundary-based shape representation, based exclusively on the edge cue.

$C|^{VP}|A$ — A consistency measure between the above A|VP|C representations and the expected models. A set of spatial deviations which might offer more successful model consistency.

$P|^{VP}|A$ — Image-based representations of edge cue derived by the peripheral holon.

$A|^{VP}|P$ — The spatial extent is trivial in this case and is confined to an elementary 3x3 neighbourhood, centred on the current attentional focus. Cue emphasis is also trivial: there is only one cue and only one resolution.

Information representations associated with self-assertive tendency

General and contour-specific.

$^{A}|EA$ — These representations are identical to the P|VP|A and C|VP|A representations of the self-assertive tendency of the holon of visual perception.

$EA|^{A}$ — Similarily, these representations are identical to the A|VP|P and A|VP|C representations.

Cognitive holon: canon of function

General statement. This holon shall analyse high-level representations of images (e.g. segmented contours) and shall check for compatibility with "previously perceived" models. It shall infer the purpose and structure of the elementary (sub-)structures, i.e., 2½-D sketch primitives. Note: this could be accomplised on the basis of a priori knowledge but it is preferable if this knowledge is not embedded but instead learned from earlier perception or experience. The holon shall hypothesise the existence of physical objects (including identification, position, and orientation) on the basis of visual data (high-level representations) and it shall seek verification of this hypothesis.

Contour-specific statement. This holon shall analyse boundary-based shape descriptors / contours (either total or partial) and shall check for compatibility with previously learned shapes. It shall infer the global structure of a shape on the basis of partial contours. Models shall be built and validated by interactive affirmation and specification of episemons. The holon shall hypothesise the actual existence of objects on the basis of available visual data (partial boundaries), including identification of episemon, position, and orientation. It shall seek verification of this hypothesis.

Cognitive holon: canon of structure

General statement. The cognitive holon is organised as two modules concerned with recognition and with inference and modelling. Recognition comprises sub-modules for analysis and association of high-level visual primitives and sub-modules for global model-based recognition and identification. The inferential and modelling sub-module shall include some canonical methodology of accessing the models and comparing them with the perceived scene on the basis of the recognition criteria of the recognition sub-modules.

Contour-specific statement. The cognitive holon is organised as two modules of recognition and of inference and modelling. Recognition comprises two sub-modules for recognition of partial contours and for recognition of global/total contours. The inferential and modelling structure shall comprise a canonical methodology for accessing models, i.e., shape prototypes, on the basis of both partial and total contours and shall compare them with perceived representations (on the basis of the recognition sub-structure), hypothesising identity and spatial configuration.

Cognitive holon: canon of integrative tendency

General statement.

A|VP|C — The cognitive holon shall base its inferences and action on the information rendered to it by the attentional holon via the holon of visual perception.

C|VP|A — It shall supply information to the attentional holon regarding the goodness of attentionally treated cues to assist the attentional holon in its operation.

A|VP|C, P|VP|C — Hypothesis checking may be effected by the cognitive holon by direct reference to cues given by the peripheral holon and higher-level image representations (e.g. 2½-D primitives) given by the attentional holon, both via the visual perception holon.

C|VP — The cognitive holon shall provide the visual perception holon with a fully identified canonical model of the environment (object-centred), together with an episematic label.

Contour-specific statement. The above general statement need only be refined by asserting the following.

C|VP|A — Contours must be consistent with expected shape; quantitative assessment will be provided regarding such consistency.

C|VP — The canonical model is a 2-D shape descriptor, fully identified (position and orientation), and including an episematic label.

Cognitive holon: canon of self-assertive tendency

General statement. The cognitive holon shall derive consistent simple inferences (formulate hypotheses) regarding the visual scene on the basis of information available to it. The operation of the holon is based on the mutually co-operative function/interaction of two sub-holons corresponding to the recognition and inferential/modelling processes. (Canons of the self-assertive tendency must address the coordination of the activity of these two holons, thus:)

C|R — The cognitive holon shall make available to the recognition holon (R) information derived through the integrative tendency, i.e., C|R ≡ A|VP|C + P|VP|C.

R|C|I/M — The recognition holon shall derive consistency measures between high-level primitives on a local and semi-local basis and on a global basis; the latter to include complete identification. These measures are available to the inferential/modelling holon.

I/M|C|R — The recognition holon shall be guided by the inferential/modelling holon when initiating recognition in alternative spatial modes.

I/M|C, C|I/M, C|R — The inferential/modelling holon, inferring purpose and structure of primitive visual structures on the basis of learned a priori models (canonically organised), shall attempt to verify hypotheses by providing the cognitive holon with information required to test it (via attentional and peripheral holons). The test is accomplished by the inferential/modelling holon using the C|I/M ≡ C|R ≡ A|VP|C + P|VP|C representations.

I/M|C — The inferential/modelling holon finally makes available canonical models to the cognitive holon.

Contour-specific statement. The above general statement need only be refined by asserting the following.

R|C|I/M — The recognition holon shall derive consistency measures on the basis of partial contours and on the basis of global contours.

I/M|C — The canonical model is a 2-D shape descriptor.

Information representations associated with integrative tendency

General and contour-specific.

$A|^{VP}|C$, $C|^{VP}|A$, $P|^{VP}|C$, $C|^{VP}|P$, and $C|^{VP}$ representations are as detailed in the information representations associated with the self-assertive tendency of the holon of visual perception.

Information representations associated with self-assertive tendency

General.

$C_{|R}$ Higher-level image-centred scene representations, typically of the form of the primal sketch, i.e. coherent semi-local 2½-D representations based on several visual cues (A|VP|C), and, in addition, image-based representations of all low-level cues derived by the peripheral processes (i.e. P|VP|C).

$R|^{R}|I/M$ Consistency measures on local and on global/emergent bases.

$I/M|^{C}|R$ Spatial deviation and canonical representation of models.

$C_{|I/M}$ $\equiv C|I/M \equiv C|R$

$I/M|^{C}$ Canonical model and re-attentional information (e.g. extent and focus).

Contour-specific.

$C_{|R}$ Image-centred boundary-based shape descriptor, in particular an NCD descriptor, based exclusively on edge cue (i.e. A|VP|C), and, in addition, image-based representation of edge cue derived by the peripheral process (i.e. P|VP|C).

$R|^{C}|I/M$ Consistency measures on local and on global/emergent bases, i.e., partial and global contours.

$I/M|^{C}|R$ Spatial deviation and canonical representation of models.

$C_{|I/M}$ $\equiv C|I/M \equiv C|R$

$I/M|^{C}$ Identified 2-D shape representation and re-attentional information (e.g. extent and focus).

Summary

Holarchies, i.e. hierarchies of holons, represent an attempt to address the problems posed by complex systems and, in particular, by biological and sociological phenomena. The approach is motivated by the inadequacy of classical reductionist mechanistic techniques and by the observation that such systems possess holistic properties. It is intended as a combination of the valid aspects of both holism and reductionism. Its keystone is the concept of the holon, a Janus-faced entity displaying integrative and self-assertive tendencies simultaneously. These holons are organised in a hierarchic fashion into holarchies: each holon represents some distinct sub-whole within the overall organisation. The holon has no complete meaning except within the context of the holarchy, requiring the other holons to lend it the necessary relief against which it must always be viewed. This embodies the holistic aspect of the approach. Nevertheless, each holon is distinct from the others and exhibits a considerable degree of autonomy and self-sufficiency. Each holon is governed by a fixed set of rules, the canons of the holon, and its behaviour is constrained by these canons. This does not limit the strategies that may be adopted by the holon in achieving any of its objectives. These canons make explicit mechanisms by which the integrative tendencies (representing the holon's partness) and the self-assertive tendencies (representing the holon's wholeness) are manifested.

The central thesis of this paper is that visual perception and robot vision may be meaningfully investigated and implemented by adhering to the holarchic approach and, indeed, that they require that such a reductionist/holistic philosophy be addressed. The holarchy is organised, essentially, as three holons of cognitive, attentional, and peripheral processes. The system is conceived as a general visual perception system and made explicit in the context of a contour-based dynamic boundary-following system. Prior attempts at utilising this perceptual organisation in a true hierarchic manner, while satisfactory in simple environments, proved insufficient when applied to non-trivial scenes, e.g., bin-picking. This inadequacy provided significant motivation for the investigation of the approach suggested in this paper and, while a definitive evaluation of its merit has yet to be accomplised, the holarchic philosophy has provided several insights into perceptual organisation. If all the answers are not yet provided, the philosophy helps one to at least ask the right questions.

Acknowledgements

The author wishes to thank Dermot Furlong of Trinity College, Dublin, for his part in many helpful discussions during the preparation of this paper.

References

1. Koestler, A. and Smythies, J.R. (Eds.) Beyond Reductionism - the Alpbach Symposium, London, 1969.
2. Koestler, A. Janus: A Summing Up, Pan Books Ltd., London, 1979.
3. Pattee, H.H. in Towards a Theoretical Biology, Waddington, C.H. (Ed.), Edinburgh, 1970.
4. von Bertalanffy, L. Problems of Life, Harper Torch Books, New York, 1952.
5. Weiss, P. in Beyond Reductionism - the Alpbach Symposium, Koestler, A. and Smythies, J.R. (Eds.), London, 1969.
6. Bennett, J.G. The Dramatic Universe, Vol. 3, Coombe Springs Press, England, 1966.
7. Martin, W.N. and Aggarwal, J.K. Survey - Dynamic Scene Analysis, Computer Graphics and Image Processing, Vol. 7, No. 3, pp.356-374. 1978.
8. Jain, R. and Haynes, S. Imprecision in Computer Vision, Computer, Vol. 15, No. 8, pp.39-48. 1982.
9. Pau, L.F. Approaches to Industrial Image Processing and their Limitations, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. PAMI-5, No. 6, pp.602-608. 1984.
10. Lee, H.C. and Fu, K.S. The GLGS Image Representation and its Application to Preliminary Segmentation and Pre-Attentive Visual Search, IEEE Computer Society Conference on Pattern Recognition and Image Processing, pp.256-261. 1981.
11. Hanson, A.R. and Riseman, E.M. Segmentation of Natural Scenes, in Computer Vision Systems, Hanson, A.R. and Riseman, E.M. (Eds.). 1978.
12. Vernon, D. A Hierarchically-Organised Robot Vision System, Proc. AI Europa, Wiesbaden, West Germany, 1985.