PERCEPTION AND COMPUTATION

Bruce M. Bennett[1] Donald D. Hoffman[2] Chetan Prakash[3]

[1] Mathematics, UC Irvine 92717 USA [2] Cognitive Science, UC Irvine 92717 USA
[3] Mathematics, CSU San Bernardino 92407 USA

Abstract

We suggest that computation is an inadequate formal foundation for the field of perception, and propose a new formal foundation – the *observer*. We propose the **observer thesis:** *For each act of perception, regardless of modality, there is an observer to perform that act of perception.* We describe formally the relationship between Turing machines and observers. We discuss conditions in which observers may be simulated by Turing machines.

1. Introduction

The investigation of specific problems in a scientific field is aided by a general theory, one that captures the formal structure underlying each specific problem, but that discards the inessential details of each problem. For instance, in the field of computer science, the theory of Turing machines and automata provides a general theory of computation. The Turing machine gives the formal structure common to all specific computations; this formalism allows one to study computation without distraction by the inessential details of specific computations. Nevertheless, this general formalism also facilitates the study of specific computations; for example, it provides general theorems on decidability and tools for the study of computational complexity. Another instance of a general theory is the Hilbert space formulation of quantum theory, which abstracts the formal structure common to various specific physical systems. One then uses the general quantum formalism to study a particular system, for example by specifying its hamiltonian.

The field of perception has as yet no comparable formal foundation. Many general principles have emerged from the detailed study of specific perceptual problems, but these principles have heretofore not been formulated into a single coherent formalism. Such a formalism would capture the structure common to all specific acts of perception, but would abstract away from the inessential details of particular perceptual problems. It would play the same role for the field of perception that the Turing machine plays for the field of computer science.

We propose such a structure, a formalism we call an *observer*. We further propose the following **observer thesis:** *For each act of perception, regardless of modality, there is an observer to perform that act of perception.* We suggest that this thesis plays a role in the field of perception analogous to the role of the Church-Turing thesis in the field of computation. We note that the class of observers properly contains the class of Turing machines: we describe a natural embedding of the class of Turing machines into the class of observers, and call this subclass of observers "Turing observers". We suspect that most psychologically plausible observers are not Turing observers; we produce an example non-Turing observer in the case of recovering the 3-D structure of moving points from dynamic images. This leads us to ask to what extent non-Turing observers can at least be simulated by Turing machines at a given scale. We examine one obvious discretization procedure, by means of which this question may be addressed. We find that there are non-Turing observers which, even after being discretized by this procedure, cannot be simulated by Turing machines. However, many observers of interest in computer vision have discretizations that are indeed Turing-simulable.

The deeper questions about the relationship of perception and computation are concerned not with a single discretization (at some given scale) but with the collection of all discretizations; indeed it is the single observer of which these are the discretizations which gives mathematical unity to the collection. Moreover it appears that there is a perceptual unity to this collection (for example, in the case of the perception of structure from motion by human observers). It seems unreasonable to assume that this unity is grounded in the computational algorithms, either associated to the individual discretizations or to their comparison. This is true in the same sense that it is unreasonable to equate the *concept* of an integral with the totality of numerical algorithms for its approximation at various scales. We suggest in this spirit that the fundamental character of perception is not computational.

CH2465-3/87/0000/0356$01.00 © 1987 IEEE

2. Definition of observer

The following definition of observer is illustrated in Figure 2.2. The mathematical notation and terminology used here are collected in the appendix.

Definition 2.1. An *observer* is a six-tuple, $((X, \mathcal{X}, \mu_X), (Y, \mathcal{Y}), E, S, \pi, \eta)$, where:

1. $(X, \mathcal{X})$ and $(Y, \mathcal{Y})$ are measurable spaces. μ_X is a measure class on X. $E \in \mathcal{X}$ and $S \in \mathcal{Y}$.

2. $\pi: X \to Y$ is a measurable surjective function with $\pi(E) = S$.

3. $\mu_X(\pi^{-1}(S) - E) = 0$.

4. Let $(E, \mathcal{E})$ and $(S, \mathcal{S})$ denote the measurable spaces on E and S respectively induced from those of X and Y. Then η is a markovian kernel on $S \times \mathcal{E}$ such that, for each s, $\eta(s, \cdot)$ is a probability measure supported in $\pi^{-1}(s) \cap E$.

By an abuse of notation, we sometimes will use μ_X to denote a representative measure of the measure class. Also we will often write X for $(X, \mathcal{X})$ and Y for $(Y, \mathcal{Y})$ when the meaning is clear from the context.

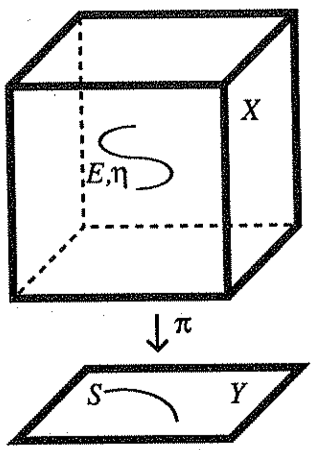

FIGURE 2.2. *An illustration of the definition of observer.*

Intuitive discussion of the definition of observer

An observer is an inferencing system, one whose conclusions are not, in general, logically implied by its premises. Roughly, the possible premises for an observer are modeled by its space Y, and the possible conclusions by its measures $\eta(s, \cdot)$ supported on E.

E is called the *configuration event* of the observer or, for brevity, the *event* of the observer. The space X is called the *configuration space* of the observer. X is a formal representation of those possible states of affairs over which the configuration event of the observer is defined. The possible perceptual conclusions of an observer are probability measures supported in the configuration event E.

The space Y is called the *observation space* of the observer. It is a formal representation of the premises available to the observer for making inferences about occurrences of E.

The measurable function π is called the *perspective map* of the observer.

The set S is called the *observation event* of the observer. All and only points in S are premises of observer inferences which conclude that an instance of the configuration event E has occurred.

The kernel η is called the *conclusion kernel* of the observer. For each $s \in S$, i.e. for each point in the observation event, $\eta(s, \cdot)$ is a probability measure on $(E, \mathcal{E})$ supported on $\pi^{-1}(s) \cap E$. Intuitively, this measure gives the relative likelihood assigned by the observer to various subsets of its event E, in consequence of the observer being given the "proximal stimulus" s. One can think of the kernel η as a convenient way of assigning to every point of S a probability measure on E. Thus η associates a perceptual conclusion to each set of premises, such that the conclusion lies in that set of points of E that can give rise to the proximal stimulus s, given the perspective map π. In this sense, η is the "rule of inference" used by the observer.

Intuitively, here is how the observer structure works in practice. When the observer is presented with an appropriate state of affairs in the world, there is some point $x \in X$ that represents that state of affairs. Then π maps x to some point y of the observation space Y. Informally, we say that the point y "lights up" for the observer. If x was a point of E, then the y that subsequently lights up is some point of S. If x was not in E and not in $\pi^{-1}(S) - E$, then the associated y will not be in S. All the observer has direct access to is which point of Y lit up, not to x. The observer wants to guess x. If the point of Y that lights up is not in S, then the observer decides that x was not in E. The observer indicates this decision by doing nothing. If the point of Y that lights up is in S, then the observer decides that x was in E, but the observer does not in general know precisely which point of E it was. In consequence, the observer couches its guess in the form of a probability distribution $\eta(s, \cdot)$ supported on E. If there is no ambiguity as to which point of E x must have been (given that it must be in E), then this probability distribution simply becomes Dirac measure supported on the appropriate point of E.

More discussion of the definition of observer

We discuss the four conditions listed in the definition of observer.

Condition 1: $(X, \mathcal{X})$, $(Y, \mathcal{Y})$ *are measurable spaces.*

μ_X*, is a measure class on* X. $E \in \mathcal{X}$ *and* $S \in \mathcal{Y}$. X is a mathematical representation of those possible states of affairs over which the configuration event E of the observer is defined. *X itself is not the real world, but a mathematical representation of some possible states of affairs in the real world.* Y is that "projection" of X from which the observer can make inferences about occurrences of E. X and Y are specified to be measurable spaces because this is the least restrictive assumption that always allows us to discuss the probabilities (more generally, measures) of events on these spaces. It would be unnecessarily restrictive to specify that X must be, say, a Euclidean space or a manifold.

μ_X is a set of measures on X that we intuitively think of as being "unbiased". This means that their definition makes no reference to properties of E or π. We may think of μ_X as expressing an abstract uniformity of X which exists prior to the notion of the configuration event set E of interest to the particular observer. For example, μ_X might be an invariant measure class for some group action on X. Again intuitively, the role of μ_X is to provide an unbiased background "probability" in terms of which the observer can be represented as an ideal decision maker (in a sense which we will discuss below), and to which the probabilities of observation of configuration events *actually* occurring in some concrete universe can be compared.

Condition 2: $\pi: X \to Y$ *is a measurable surjective function with* $\pi(E) = S$. π must be surjective for otherwise there would be points in the observation space, Y, that did not arise from the configuration space, X. That is, the observer would have observation points that were gratuitous, and then we could simply discard them while preserving all the other hypotheses. π must be measurable, for in this theory the statistical information available at the sensorium Y must at the very least be semantically compatible with information about probabilities of states of affairs in the world as represented in configuration space X.

Condition 3: $\mu_X(\pi^{-1}(S) - E) = 0$. This condition means that the probability of "false targets" is zero. Intuitively, a false target is a distal stimulus that leads an observer to an illusory perception. More precisely, a false target for an observer is an element of $\pi^{-1}(S) - E$.

Because of condition 3 an observer is an ideal decision maker in the following sense: *Given that the state of affairs in X is not in E, the observer almost surely recognizes that the state of affairs is not in E.* Put another way, the measure of false targets is zero for an unbiased measure μ_X on the configuration space X. This is because points in X that are not in E almost surely do not project via π to S and are therefore not given an interpretation by the observer.

It is also true that *Given that some event in E has occurred, the observer always recognizes this.* Occurrences of events in E always lead the observer to select as its conclusion a probability measure supported on E, simply because $\pi(E) = S$.

State of Affairs

Decision	E	-E
E	1 "Hit"	0 "False Alarm"
-E	0 "Miss"	1 "Correct Reject"

FIGURE 2.3. *Decision diagram for observers.*

The sense in which an observer is an ideal decision maker can also be represented in a decision diagram, as shown in Figure 2.3. The diagram represents two possible states of affairs across the top: E, which indicates that some instance of the observer's configuration event has occurred, and -E, which indicates that an event other than the configuration event has occurred. The diagram represents the two possible decisions of the observer along the side. Inside each box in the right column is a number which is a conditional probability, namely the unbiased conditional probability that the observer arrives at the decision indicated to the left side of the diagram given that E does not occur. Inside each box in the left column is a number; in this left column the number 1 is a shorthand for "certainly" and 0 for "certainly not". The numbers in this left column hold simply by the definition of observer; if a configuration in E occurs, then since $S = \pi(E)$ and the observer always decides that E occurred given a premise in S, the observer always decides correctly. Also inside each box is a label in quotes which describes the type of decision represented by that box. So, for instance, the box labelled "false alarm" has the number 0 in it. This means that, for an observer as specified in Definition 2.1, the conditional probability is zero that the observer will decide that an event in E occurred given that in fact an event outside of E occurred. (The one in the box labelled "correct reject" is the complementary conditional probability).

An observer is an ideal decision maker in this sense *regardless of the relationship between the observer and the world it is in.* However, for most observers, without the proper relationship between the observer and the world, the observer's decision that E occurred will often be incorrect. This is because for most observers there will be points in $\pi^{-1}(S) - E$, and the unbiased measure μ_X may, in assigning zero measure to this set, be contradicting actual states of affairs in the world. Note that the order of the probabilistic conditioning is crucial to a clear understanding of this situation.

That aspect of the observer inference presented in the decision diagram of Figure 2.3 is not the only one of interest. The observer does not merely decide whether or not an event in E has occurred: it selects a probability measure supported on E which is the observer's best guess as to which events in E are likely to have occurred, together with their likelihoods. One would like to know if the particular probability measure selected by the observer is accurate. This question involves establishing a formal framework in which observer and observed can be discussed (we do this in Bennett et al. (1987)). The question of perceptual accuracy can then be understood in terms of stabilities of dynamics of participators on these frameworks. In particular, we can ask whether the conclusion kernel η of the observer is compatible with these stabilities; this leads to "perception=reality" equations.

Condition 4: *η is a markovian kernel on $S \times \mathcal{E}$ such that, for each s, $\eta(s, \cdot)$ is a probability measure supported on $\pi^{-1}(s) \cap E$.* η captures formally the conclusions made by an observer for premises in S. A conditional probability distribution is a natural formal object for this purpose because for each $s \in S$ it assigns a probability measure whose support is a subset of the fibre over s under the map π. That is, the only points $e \in E$ that are involved in the perceptual conclusion for the premises s are those such that $\pi(e) = s$.

A remark on noise

Observer theory replaces noise with perceptual uncertainty. By perceptual uncertainty we mean firstly that the map $\pi|_E$ is, in general, many to one, and in particular the conclusion measures $\eta(s, \cdot)$ of an observer are supported, in general, on sets of more than one element in the fibres of π. Thus there is fundamental uncertainty about which configuration event corresponds to a given observation event. Secondly, in general E intersects the fibres of π in proper subsets. Thus there is fundamental uncertainty here as to whether a particular observation in S resulted from a configuration event in E.

The notion of noise entails the idea of an independent objective reality grounded in a fixed framework: an object exists in spacetime in some precise state but the measurement of that state is "interfered with" by imprecision of measurement. Even in quantum mechanics one can maintain this objectivist view of noise; the imprecision is there ordained by principle (Heisenberg's). We take a completely different tack. Perception is the primary phenomenon, and it has a fundamentally probabilistic character which expresses itself in distributions on the spaces of premises and conclusions. We assume, moreover, that each premise may be viewed (possibly in a nonunique way) as a set of conclusions of a system of "lower level observers", which we call an instantiation of the given observer. It follows that the distributions of conclusions at the level of the instantiation become the premises at the next level. In this manner, we view perceptual uncertainty as propagating up a lattice of observers, and so it is not grounded in the imprecision of measurement of some objectively existing system in some fixed spacetime.

3. An example

We present an example of an observer, one which infers the three-dimensional structure of rigid objects that spin rigidly about a fixed axis. We construct this "structure from motion" observer based on the following result (Hoffman and Bennett, 1986):

- (i) Given three distinct orthographic projections of three points in $\mathbf{R}^3$ that are spinning rigidly about a fixed axis, the 3-D structure and motion of the points is almost surely determined uniquely (up to a reflection about a plane parallel to the image plane). Moreover (ii) the Lebesgue measure is zero of the set of those image data that permit such a determination.

Because of this result we can construct an observer that infers the three-dimensional structure of three points in rigid fixed-axis motion from three orthographic views. By "three-dimensional structure" of the three points we mean their positions relative to each other. The observer takes one of the points to be the origin O and represents the positions of the other two points A_1 and A_2 relative to that origin. In this case the configuration space X is the space of all three-tuples of pairs of points, where each point lies in $\mathbf{R}^3$. That is,

$$X = \{(\mathbf{a}_{ij}) \,|\, \mathbf{a}_{ij} = (x_{ij}, y_{ij}, z_{ij}); i = 1,2; j = 1,2,3\} = \mathbf{R}^{18}.$$

The observation space Y is the set of all triples of pairs of points in $\mathbf{R}^2$, i.e.

$$Y = \{(\mathbf{b}_{ij}) \,|\, \mathbf{b}_{ij} = (x_{ij}, y_{ij}); i = 1,2; j = 1,2,3\} = \mathbf{R}^{12}.$$

The perspective map is then $\pi: \mathbf{R}^{18} \rightarrow \mathbf{R}^{12}$ induced by $(x_{ij}, y_{ij}, z_{ij}) \mapsto (x_{ij}, y_{ij})$. The unbiased measures μ_X and μ_Y can be taken to be Lebesgue measure. The σ-algebras $\mathcal{X}$ and $\mathcal{Y}$ are the appropriate Borel algebras.

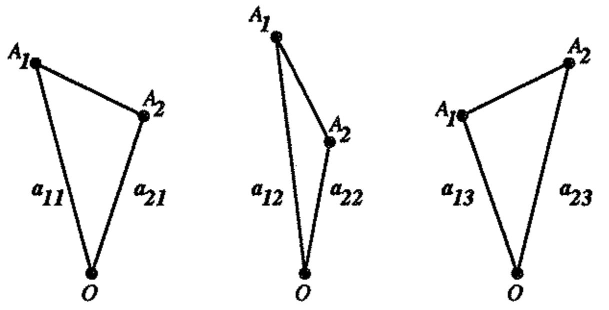

FIGURE 3.1. *Rigid fixed-axis motion: Three views of three points*

To define the configuration event E we use notation as illustrated in Figure 3.1. The three points are O, A_1, and A_2. As above, let $\mathbf{a}_{ij}$ denote the vector in three dimensions between points O and A_i in view j ($j = 1, 2, 3$). E is that subset of X consisting of three pairs of points, each point of the pair lying in $\mathbf{R}^3$, such that there is a rigid translation and rigid rotation about a single axis relating each pair plus the origin point to the others. It happens in this case that E is an algebraic variety (the solution set of polynomial equations) defined by the following eight vector equations:

$$\mathbf{a}_{11}\cdot\mathbf{a}_{11} - \mathbf{a}_{12}\cdot\mathbf{a}_{12} = 0, \quad (3.2)$$
$$\mathbf{a}_{11}\cdot\mathbf{a}_{11} - \mathbf{a}_{13}\cdot\mathbf{a}_{13} = 0, \quad (3.3)$$
$$\mathbf{a}_{21}\cdot\mathbf{a}_{21} - \mathbf{a}_{22}\cdot\mathbf{a}_{22} = 0, \quad (3.4)$$
$$\mathbf{a}_{21}\cdot\mathbf{a}_{21} - \mathbf{a}_{23}\cdot\mathbf{a}_{23} = 0, \quad (3.5)$$
$$\mathbf{a}_{11}\cdot\mathbf{a}_{21} - \mathbf{a}_{12}\cdot\mathbf{a}_{22} = 0, \quad (3.6)$$
$$\mathbf{a}_{11}\cdot\mathbf{a}_{21} - \mathbf{a}_{13}\cdot\mathbf{a}_{23} = 0, \quad (3.7)$$
$$(\mathbf{a}_{11} - \mathbf{a}_{12})\cdot[(\mathbf{a}_{11} - \mathbf{a}_{13})\times(\mathbf{a}_{21} - \mathbf{a}_{22})] = 0, \quad (3.8)$$
$$(\mathbf{a}_{11} - \mathbf{a}_{12})\cdot[(\mathbf{a}_{11} - \mathbf{a}_{13})\times(\mathbf{a}_{21} - \mathbf{a}_{23})] = 0. \quad (3.9)$$

In these equations the operation $\cdot$ indicates scalar (dot) product and $\times$ indicates vector (cross) product. The first six equations specify that the three points move rigidly. The last two specify that the points rotate about a fixed axis. E so defined has positive codimension in X (i.e. the dimension of E is less than that of X) and so E has Lebesgue measure zero in X. The observation event is $S = \pi(E)$, and has Lebesgue measure zero in Y. Therefore the Lebesgue measure of false targets is zero (i.e. $\mu_X(\pi^{-1}(S) - E) = 0$). With effort it can be shown that generically on S the fibre of π over a point $s \in S$, i.e. the set of points $x \in X$ such that $\pi(x) = s$, contains two points of E (Hoffman and Bennett, 1986). We can chose η to be the conditional probability distribution on E relative to π which gives weight, say, of one half to each of the two points. Abstractly, this observer structure is as follows:

$$\begin{array}{ccc} X = \mathbf{R}^{18} & \supset & E \quad (= \text{rigid fixed-axis motions}) \\ \downarrow \pi & & \downarrow \pi \\ Y = \mathbf{R}^{12} & \supset & S \end{array} \quad (3.10)$$

4. Turing observers

All Turing machines have sufficient structure to be viewed as observers. The set of nontrivial Turing machines is a small proper subset of the set of observers; intuitively, it is a subset of observers whose inferences are deductively valid. Observers more generally perform inferences that are not deductively valid, but that have some degree of inductive strength. Moreover, most observers are not Turing observers if only because the sets E and S need not be, in general, recursively enumerable. (Recursively enumerable sets are precisely those that can be recognized by Turing machines.)

The theory of automata considers several characterizations of Turing machines. All characterizations are equivalent to defining a Turing machine as a language recognizer. For example, one can characterize a Turing machine as computing a partial recursive function f; in this case the graph of f is a recursively enumerable set, and computing f is equivalent to recognizing its graph. Therefore we will view Turing machines as recognizers of recursively enumerable languages. Let Σ be the terminal alphabet of a Turing machine T, Σ^* be the set of all strings of finite length over the alphabet Σ, and $L \subset \Sigma^*$ be the language recognized by T. We take for the σ-algebra of Σ^* simply its power set. We then define an embedding of the set of Turing machines into the set of observers by the map

$$T \mapsto (\Sigma^*, \Sigma^*, L, L, \text{identity}, \epsilon)$$

where, if $l \in L$, and ϵ_l denotes Dirac measure concentrated at l, then ϵ is given by $\epsilon(l, \cdot) = \epsilon_l$. Note that $\pi^{-1}(S) = E$ so that for any measure class on Σ^*, $(\Sigma^*, \Sigma^*, L, L, \text{identity}, \epsilon)$ is an observer. Moreover, since π is the identity, so that the fibre $\pi^{-1}\{s\}$ equals $\{s\}$ for each point $s \in S$, $\eta(s, \cdot)$ is Dirac measure concentrated on this point. It is straightforward to verify that, with these identifications, all the restrictions placed on observers by Definition 2.1 are satisfied.

Once we realize the Turing machines as a subclass of observers in this way, it is obvious that most observers are not Turing machines. More generally, however, we ask whether a Turing *simulation* exists for a given observer. For this purpose we would like to define a *canonical simulation* procedure. In fact, we will discuss the simulation of structures slightly more general than observers, namely we will drop the measure zero requirement in condition 3, since no property of the measure class μ_X will be relevant to the simulation definition. We will call these more

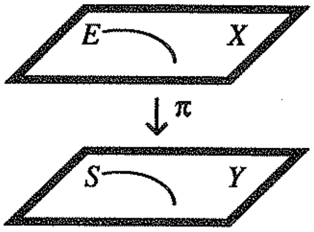

FIGURE 4.1. *A Turing observer.* $X = Y = \Sigma^*$. $E = S = L$. *π being an isomorphism means that the Turing observer's conclusions are deductively valid.*

general structures *quasi-observers*, still using the notation $O = (X, Y, E, S, \pi, \eta)$ for a quasi-observer. Let $O = (X, Y, E, S, \pi, \eta)$ be the quasi-observer to be simulated. The objective of the simulation is the computation of $\eta(s, A)$, for all relevant sensorial points s and configurational situations A. Now clearly such a computation is possible only if the "relevant" sets A can be seen as generated by a countable collection of events. To this end, let $\mathcal{A}$ denote the set of atoms of the σ-algebra $\mathcal{X}$ of X. By definition, the elements of $\mathcal{A}$ are the minimal measurable sets A of $\mathcal{X}$, i.e. no nonempty proper subset of A is measurable. Similarly we call a finite real-valued measure μ on X *atomic* if the minimal sets of nonzero μ-measure are atoms of $\mathcal{X}$ (that is, if whenever $B \in \mathcal{X}$ with $\mu(B) \neq 0$, then B contains an atom A of $\mathcal{A}$ such that $\mu(A) \neq 0$.) Any such measure induces a function from $\mathcal{A}$ to $\mathbb{R}$, viz. $A \mapsto \rho(A)$; indeed, this function is informationally equivalent to the atomic measure. In particular, for a fixed $s \in S$, $A \mapsto \eta(s, A)$ is a real-valued function on $\mathcal{A}$. We require, for purposes of simulation, that the measures $\eta(s, \cdot)$ be atomic (for otherwise our scheme below may simulate a vacuous portion of the quasi-observer). Thus, we may associate to O the function $f: S \times \mathcal{A} \to \mathbb{R}$ defined as follows:

$$f(s, A) = \eta(s, A).$$

Assuming that the measures $\eta(s, \cdot)$ are atomic, we define the *canonical Turing simulator of* O to be the machine T which recognizes S in Y and then computes f. Thus T exists if and only if O satisfies the requirements:

i. S is recursively enumerable in Y.
ii. f is recursive.

For most observers of interest to vision researchers, $\eta(s, \cdot)$ is not atomic, and S is uncountable (so that f has little chance of being recursive). Hence for these observers there is no canonical simulator. For example, in the observer presented in section 3 the measures $\eta(s, \cdot)$ are atomic (almost surely on S) since generically there are only two points of E over each point of S. However, since both S and E are uncountable the function f will not be recursive, and there is no canonical simulator. Even when everything is countable, f will not be recursive in general, so simply by making a discrete approximation to an observer we cannot expect that it will have a Turing simulation. However, at least in certain instances of interest to vision researchers, discrete approximations may allow Turing simulation. For these reasons and others it is essential to have a general theory of discretization of observer structures. We now give some indication of this.

For our present purposes, we will restrict attention to Euclidean configuration spaces X and observation spaces Y (with their Borel algebras), and assume that $\pi: X \to Y$ is projection. Thus, let $O = (X, Y, E, S, \pi, \eta)$ be an observer with $X = \mathbb{R}^{n+m}$, $Y = \mathbb{R}^n$, and π projection, say onto the first n coordinates. In order to effect a discretization an additional datum is required, namely a measure λ on S. Intuitively, λ and η come from the same source, namely a probability measure ρ on E which expresses the actual occurrence probabilities of configuration events in a specific universe. In particular, in this case the natural choice for λ is $\pi_*(\rho)$, just as the natural choice for η is a version of the conditional probability distribution of ρ with respect to π. In our case, since λ and η are assumed given, we can simply define the measure ρ on E by $\rho = \lambda\eta$, i.e.

$$\rho(A) = \int_S \lambda(ds)\eta(s, A), \quad A \in \mathcal{E}.$$

We will describe the resulting canonical discretization procedure in terms of this ρ. This procedure will result in quasi-observers with countable configuration spaces. Indeed this is the motivation for considering the more general class of quasi-observers. It is unreasonable to expect that any unbiased measure μ on a countable space would satisfy the measure zero requirement in condition 2 of the observer definition.

For each rational $\delta > 0$, we can partition X and Y into hypercubes whose edges have length δ, and whose vertices have coordinates which are integer multiples of δ. Let us call the resulting sets of hypercubes X_δ and Y_δ. π induces a map $\pi_\delta: X_\delta \to Y_\delta$. Let E_δ denote those hypercubes $\bar{x}$ of X_δ such that $\rho(E \cap \bar{x}) > 0$. Let $S_\delta = \pi_\delta(E_\delta)$. As a consequence of these definitions, if $\bar{e} \in E_\delta$ then $\rho(\bar{e}) > 0$, and if $\bar{s} \in S_\delta$, $\lambda(\bar{s}) > 0$. We will define below a kernel η_δ (depending on the original kernel η and δ) such that $O_\delta = (X_\delta, Y_\delta, E_\delta, S_\delta, \pi_\delta, \eta_\delta)$ is a quasi-observer. We can think of this quasi-observer O_δ as a "δ-discretization" of O.

This, however, is not sufficient for our purposes. In fact, we we would like to be able to compare the various discretizations (at different scales δ) with each other and the original observer. To this end, we seek a canonical way of embedding the discrete spaces E_δ and S_δ as subsets E'_δ and S'_δ of the original X and Y, *in such a way that the original perspective map π is retained as the perspective map of the new quasi-observer* $O'_\delta = (X, Y, E'_\delta, S'_\delta, \pi, \eta_\delta)$.

To achieve this let us first consider how to embed S_δ in Y. Given a hypercube $\bar{s} \in S_\delta$, we may find its center of mass with respect to the restriction of the measure λ to $\bar{s}$. This center of mass will not, in general, lie in S, but it is the natural punctual representative of $\bar{s}$ in Y. Recalling that for $\bar{s} \in S_\delta$, $\lambda(\bar{s}) > 0$, we may now define the embedding $\beta: S_\delta \to Y$ by

$$\beta(\bar{s}) = \int_{\bar{s}} s \, \lambda_{\bar{s}}(ds)$$

with

$$\lambda_{\bar{s}}(ds) = \frac{1}{\lambda(\bar{s})} 1_{\bar{s}}(s)\lambda(ds), \quad \bar{s} \in S_\delta.$$

That is, $\lambda_{\bar{s}}$ is the normalized restriction of λ to the hypercube $\bar{s}$.

Similarly, we wish to define a center-of-mass embedding for E_δ using appropriate measures on X. For purposes of finding the center of mass of $\bar{e}$ in E_δ, it may seem natural to use the normalization of the restriction of ρ to $\bar{e}$. However, as we shall see below, a slightly different choice of measure on $\bar{e}$ is much better suited to the task at hand. To this end, let $\rho_{\bar{e}}$ be the normalized restriction of ρ to $\bar{e}$, that is,

$$\rho_{\bar{e}}(C) = \frac{\rho(C \cap \bar{e})}{\rho(\bar{e})}.$$

By construction of E_δ, this yields a probability measure on $\bar{e}$. It is straightforward to verify that since η is the regular conditional probability distribution (rcpd) of ρ with respect to π, the measure $\rho_{\bar{e}}$ also possesses an rcpd with respect to π, a version of which is given by the formula

$$\eta_{\bar{e}}(s, de) = \frac{\eta(s, de)}{\eta(s, \bar{e})} 1_{\bar{e}}(e)$$

(which is defined up to a set of $\pi_* \rho_{\bar{e}}$-measure zero in its first argument, and which we may take to be a markovian kernel off this zero-measure set). As usual, by composing $\eta_{\bar{e}}$ with the measure $\pi_* \rho_{\bar{e}}$ we can reconstruct $\rho_{\bar{e}}$. We shall, however, compose $\eta_{\bar{e}}$ with $\lambda_{\pi(\bar{e})}$ instead, defining a new measure $\nu_{\bar{e}}$ as

$$\nu_{\bar{e}}(C) = \int_{\pi(\bar{e})} \lambda_{\pi(\bar{e})}(ds) \cdot \eta_{\bar{e}}(s, C), \quad \bar{e} \in E_\delta,$$

where C is any measurable subset of $\bar{e}$. This is, by construction, a probability measure supported on $\bar{e}$, which gives the embedding of E_δ in X by the map α as follows:

$$\alpha(\bar{e}) = \int_{\bar{e}} e \, \nu_{\bar{e}}(de), \quad \bar{e} \in E_\delta.$$

We will denote the image $\alpha(E_\delta)$ in X by E'_δ and the image $\beta(S_\delta)$ in Y by S'_δ.

We now show that these embeddings α and β respect the original map π in the following precise sense: the center of mass of $\bar{e}$ projects to the center of mass of $\bar{s} = \pi_\delta(\bar{e})$, i.e., for $\bar{e} \in E_\delta$, $\pi(\alpha(\bar{e})) = \beta(\pi_\delta(\bar{e}))$. This is satisfying, as it allows us to use a consistent perspective map at all levels of scale δ, a fact which expresses a unification of the discretizations at the various scales in a manner which has at least a chance of being Turing computable in special cases of interest.

To see why $\pi(\alpha(\bar{e}))$ should equal $\beta(\pi_\delta(\bar{e}))$, note that since π is linear, we may take π inside the integral defining $\alpha(\bar{e})$, so that

$$\begin{aligned}\pi(\alpha(\bar{e})) &= \int_{\bar{e}} \pi(e) \nu_{\bar{e}}(de) \\ &= \int_{\pi(\bar{e})} \lambda_{\pi(\bar{e})}(ds) \int_{\bar{e}} \eta_{\bar{e}}(s, de) \pi(e).\end{aligned}$$

But $\eta_{\bar{e}}(s, \cdot)$ is supported on the fibre where $\pi(e) = s$, so that

$$\begin{aligned}\pi(\alpha(\bar{e})) &= \int_{\pi(\bar{e})} \lambda_{\pi(\bar{e})}(ds) \cdot s \int_{\bar{e}} \eta_{\bar{e}}(s, de) \\ &= \beta(\pi_\delta(\bar{e})).\end{aligned}$$

It is worth pointing out here, that had we used the full measure $\rho_{\bar{e}}$ in defining the embedding α of E_δ, this result would not have obtained.

It remains to discretize η, i.e. we want to define η_δ, so that we can speak of the quasi-observers $O_\delta = (X_\delta, Y_\delta, E_\delta, S_\delta, \pi_\delta, \eta_\delta)$, or $O'_\delta = (X, Y, E'_\delta, S'_\delta, \pi, \eta_\delta)$. The appropriate discretization of η is then the kernel η_δ given by

$$\eta_\delta(\bar{s}, \{\bar{e}\}) = \int_{\bar{s}} \lambda_{\bar{s}}(dt) \eta(t, \bar{e}), \quad \bar{s} \in S_\delta, \bar{e} \in E_\delta.$$

This is by construction a markovian kernel on $S_\delta \times \mathcal{E}_\delta$. Here we are merely averaging the contributions from the various fibres in the appropriate hypercubes. Note that we can view η_δ as a kernel on $S'_\delta \times \mathcal{E}'_\delta$ if we wish, simply by using the identifications α and β.

In general, E_δ and S_δ need not be recursively enumerable, and a fortiori the function f defined as above using η_δ need not be recursive. Thus a discretization of a non-Turing observer may not have a Turing simulation.

We now discuss the discretization procedure applied to observers of the type of the structure-from-motion observer O of §3. In that particular case (with the notation of §3) E is the locus of points in $\mathbf{R}^{18}$ satisfying the equations (3.2)–(3.9), and S is the image of E in $\mathbf{R}^2$ by the projection π. Note that the polynomial equations defining E have integer coefficients. Thus we can apply the following general result:

Suppose $Y = \mathbf{R}^r$, $X = \mathbf{R}^{r+n}$, and $\pi: X \to Y$ is projection onto a set of r of the coordinates of X. Suppose E is the locus of zeroes in X of a finite set Σ of polynomial equations (in the $r+n$ variables of X) with integer coefficients. Let $S = \pi(E)$, let δ be a rational number, and let X_δ, Y_δ, E_δ, S_δ, π_δ be as defined in the discretization procedure above. Then

i. S_δ is a recursively enumerable subset of Y_δ.

ii. For all $\bar{y} \in Y_\delta$, $\pi^{-1}(\bar{y}) \cap E_\delta$ is a recursive subset of $\pi_\delta^{-1}(\bar{y})$ (and of X_δ).

This result is obtained by applying the Theorem of Tarski on the decidability of polynomial inequalities (see, e.g. Jacobson 1974); we omit the details here.

Condition (i) of the above result corresponds to the first requirement for the Turing simulator of O_δ to exist. Condition (ii) is a necessary condition for the function f associated to the observer O_δ to be recursive, but it is certainly not sufficient for this purpose; the real issue here is the nature of η_δ. For example, even if there are only finitely many points in $\pi^{-1}(\bar{s} \cap E_\delta)$ for each $\bar{s}$, and $\eta_\delta(\bar{s}, \cdot)$ assigns equal probability to each of these points, condition (ii) does not by itself imply that f is recursive.

In discussing the issue of the computability of f, we must first of all recognize that even if the original observer O had a very simple structure, η_δ may be more complex in a certain sense than the original η. For example, suppose exactly n points of E lie over each point of S, and that η assigns equal probability to each of these points. In well-behaved situations, like the algebraic examples at hand, the map $E_\delta \to S_\delta$ induced by π_δ will still be finite-to-one. However, the number of points of E_δ mapping to a given $\bar{s}$ will vary considerably with $\bar{s}$ (but will always be at least equal to n for small enough δ). This variation corresponds to the fact that the mean slope of E over S within $\bar{s}$ will vary with $\bar{s}$ in general. This raises the question of whether there is a canonical choice for the shapes of the regions in the partition of X and Y leading to the discretization. For example, in our case $X = \mathbf{R}^n$, if we make a globally defined and well-behaved coordinate change, the hypercubes defined in terms of the new coordinate system will have in general a different shape from the original hypercubes. Can we choose a coordinate change so that the discretization resulting from the hypercube decomposition in this system will yield, for sufficiently small δ, a map $E_\delta \to S_\delta$ of minimal complexity, for example, whose degree is the same as that of the original $E \to S$?

The rigorous analysis of this question is beyond the scope of this paper. However, we can consider at top level the consequences of the existence of such a coordinate system by assuming, for the sake of discussion, that our original coordinate system has the desired property. Moreover, to fix ideas, assume that the original map $E \to S$ is n to one. Thus we are assuming that for all sufficiently small δ, the maps $E_\delta \to S_\delta$ are all n to one. Notice that this is a *stability* result for the discretizations: it says that a fundamental structural property of the discretization is eventually stable as $\delta \to 0$, and in fact it stabilizes in such a way that the corresponding structural property of the limiting object (the observer O) is attained at a *finite* stage in the discretization process.

In this way, the *stable structure* of the collection of discretizations represents the structure of the original observer. Not only the stabilities of the maps $E_\delta \to S_\delta$, but also those of the η_δ as $\delta \to 0$ should in principle be included in an analysis of this kind. In any event, these considerations will lead to notions of **effective simulation** of an observer O: a system of successively finer discretizations O_δ, converging to O, whose stable structural properties (i.e., properties which hold for all sufficiently small δ) reflect the perceptually relevant properties of the original O. Thus, the fundamental structure of O is accessible at finite stages of discretization, in a manner which is independent of scale, at least for sufficiently small scales. It seems clear that in the absence of this kind of stability the Turing simulability of the individual O_δ's is an insufficient hypothesis to justify a "perception as computation" viewpoint.

Once the notion of effective simulation is given a precise definition (which we have not done here) note that its existence is a property of O and not of any single discretization of O. The analysis of O to determine whether it is effectively simulable in this sense, is an appropriate level on which to address the real issues a propos of the relationship between perception and computation.

Acknowledgements

We thank P. Eklof, L. Narens, W. Richards, R. Shore, T. Slaman for helpful discussions. This work was supported in part by National Science Foundation grant IST-8413560 and by Office of Naval Research contract N00014-85-K-0529 to D. Hoffman.

References

Bennett, B., Hoffman, D. and Prakash, C. (1987). *Observer Mechanics*, Cambridge, Massachusetts: MIT Press (forthcoming).

Hoffman, D. and Bennett, B. (1986). "The computation of structure from fixed-axis motion: rigid structures," *Biological Cybernetics*, 54, 71-83.

Jacobson, N. (1974). *Basic Algebra I*, San Francisco: Freeman.

Appendix

The definition of observer given in this paper makes use of several mathematical concepts from probability and measure theory. In this appendix we collect basic terminology and notation from these fields for the convenience of the reader.

Let X be an arbitrary abstract space, namely a nonempty set of elements called "points". Points are often denoted generically by x. A collection $\mathcal{X}$ of subsets of X is called a *σ-algebra* if it contains X itself and is closed under the set operations of complementation and countable union (and is therefore closed under countable intersection as well). The pair $(X, \mathcal{X})$ is called a *measurable space* and any set A in $\mathcal{X}$ is called a *measurable set*. In the case that $X = \mathbf{R}^n$, the smallest σ-algebra $\mathcal{X}$ containing all open balls is called the *Borel algebra* of $\mathbf{R}^n$. A map π from a measurable space $(X, \mathcal{X})$ to another measurable space $(Y, \mathcal{Y})$, $\pi: X \to Y$, is said to be *measurable* if $\pi^{-1}(A)$ is in $\mathcal{X}$ for each A in $\mathcal{Y}$; this is indicated by writing $\pi \in \mathcal{X}/\mathcal{Y}$. In this case the set $\sigma(\pi) = \{\pi^{-1}(A) | A \in \mathcal{Y}\}$ is a subσ-algebra of $\mathcal{X}$, called the *σ-algebra of π*. The simplest such real-valued function is the so-called *indicator function* 1_A of a subset $A \in \mathcal{X}$, defined by $1_A(x) = 0$ if $x \notin A$ and 1 if $x \in A$.

A *measure* on the measurable space $(X, \mathcal{X})$ is a map μ from $\mathcal{X}$ to $\mathbf{R} \cup \{\infty\}$, such that the measure of a countable union of disjoints sets in $\mathcal{X}$ is the sum of their individual measures. A property is said to hold μ almost surely (abbreviated μ a.s.) or μ almost everywhere (μ a.e.) if it holds everywhere except at most on a set of μ measure zero. A *support* of a measure is any measurable set with the property that its complement has measure zero. A *probability measure* is a measure μ whose range is the closed interval $[0, 1]$ and that satisfies $\mu(X) = 1$. A *Dirac measure* is a measure supported on a single point. If ν and μ are two measures defined on the same measurable spaces, we say that ν is *absolutely continuous with respect to* μ (written $\nu \ll \mu$) on a measurable set E if $\nu(A) = 0$ for every $A \subset E$ with $\mu(A) = 0$. A *measure class* on $(X, \mathcal{X})$ is an equivalence class of measures on $(X, \mathcal{X})$; the equivalence relation is that of mutual absolute continuity. Given a measure space $(X, \mathcal{X}, \mu)$ and a mapping p from $(X, \mathcal{X}, \mu)$ to a measurable space $(Y, \mathcal{Y})$, one can induce a measure $p_*\mu$ on $(Y, \mathcal{Y})$ by $(p_*\mu)(A) = \mu(p^{-1}(A))$. Then $p_*\mu$ is called the *distribution of p with respect to μ*, or the *projection of μ by p*.

Let $(X, \mathcal{X})$, $(Y, \mathcal{Y})$ be measurable spaces. A *kernel on X relative to Y*, or a *conditional distribution on X relative to Y*, is a mapping $N: Y \times \mathcal{X} \to \mathbf{R} \cup \{\infty\}$, such that

i. for every y in Y, the mapping $A \to N(y, A)$ is a measure on $\mathcal{X}$ which will often be denoted by $N(y, \cdot)$;
ii. for every A in $\mathcal{X}$, the mapping $y \to N(y, A)$ is a measurable function on Y which will often be denoted by $N(\cdot, A)$. N is called *positive* if its range is in $[0, \infty]$ and *markovian* if it is positive and, for all $y \in Y$, $N(y, X) = 1$. If $X = Y$ we say that N is a *kernel on X*.

If μ is a measure on Y and N is a kernel on $Y \times \mathcal{X}$, we may define a measure on X denoted by μN and given by the formula $\mu N(B) = \int_Y \mu(dy) N(y, B)$, $B \in \mathcal{X}$.

Suppose now that $\pi: X \to Y$ is a measurable mapping, and that ρ is a probability measure on X. We say that a kernel K on $Y \times \mathcal{X}$ is a version of the *regular conditional probability distribution of ρ with respect to π* (rcpd of ρ with respect to π) if

i. For $\pi_*\rho$-almost every $y \in Y$, $K(y, \cdot)$ is supported on the fibre $\pi^{-1}\{y\}$, and $K(y, \pi^{-1}\{y\}) = 1$.
ii. For every $A \in \mathcal{X}$, $\rho(A) = \int_Y (\pi_*\rho)(dy) K(y, A)$, $A \in \mathcal{X}$. That is, $\rho = (\pi_*\rho)K$.