King's Research Portal

DOI:

[10.1093/bjps/axr050](https://doi.org/10.1093/bjps/axr050)

Document Version

Publisher's PDF, also known as Version of record

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Shea, N. (2013). Inherited Representations are Read in Development. *BRITISH JOURNAL FOR THE PHILOSOPHY OF SCIENCE*, 64(1), 1-31. <https://doi.org/10.1093/bjps/axr050>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Inherited Representations are Read in Development

Nicholas Shea

ABSTRACT

Recent theoretical work has identified a tightly constrained sense in which genes carry representational content. Representational properties of the genome are founded in the transmission of DNA over phylogenetic time and its role in natural selection. However, genetic representation is not just relevant to questions of selection and evolution. This article goes beyond existing treatments and argues for the heterodox view that information generated by a process of selection over phylogenetic time can be read in ontogenetic time, in the course of individual development. Recent results in evolutionary biology, drawn both from modelling work, and from experimental and observational data, support a role for genetic representation in explaining individual ontogeny: both genetic representations and environmental information are read by the mechanisms of development, in an individual, so as to lead to adaptive phenotypes. Furthermore, in some cases there appears to have been selection between individuals that rely to different degrees on the two sources of information. Thus, the theory of representation in inheritance systems like the genome is much more than just a coherent reconstruction of information talk in biology. Genetic representation is a property with considerable explanatory utility.

- 1 *Introduction*
- 2 *Inherited Representations*
- 3 *Reading Genetic Representations*
 - 3.1 *Do genes carry correlational information?*
- 4 *Selection Between Genetic and Environmental Information*
 - 4.1 *Modelling*
 - 4.2 *Empirical applications*
 - 4.3 *Maternal effects*
- 5 *Genetic Representation and the Genome*
 - 5.1 *Information capacity of organisms' genomes*
 - 5.2 *Many amino acids, few nucleotides*
 - 5.3 *A function of sex*

6 *Explaining Further Aspects of Development*

6.1 *Canalization against environmental variation*

6.2 *An informational function for the nuclear membrane?*

7 *Conclusion*

1 Introduction

Genetic information is controversial. Often genetic differences correlate with phenotypic differences and thereby carry information in the correlational sense analysed by information theory (Shannon information, Kullback–Leibler information, etc.). But such correlational information is ubiquitous. Shea ([2007a]) and Bergstrom and Rosvall ([2010]) have recently argued that DNA also carries information in a stronger sense, on the basis that it has acquired the adaptive function of transmitting information down the generations (see also Godfrey-Smith [2010]; Shea [2011a]). That function can form the basis of a teleosemantic theory according to which DNA represents whole organism phenotypes.¹

Sceptics about the importance of genetic representation sometimes concede that it has an explanatory role in organizing facts about evolution, genetics and selection, but argue that it can have no role to play in explaining development (Lehrman [1970], p. 136). This article answers that challenge, showing that some aspects of individual development can be explained in terms of the representational contents carried by genes.

It is uncontroversial to observe that organisms sometimes achieve an adaptive fit to a feature of their environment by detecting information about that feature as they develop. It is highly controversial to claim that organisms sometimes achieve an adaptive fit to a feature of their environment by reading information about that feature from their genes. Theorists have been keen to reject the existence of genetic representation precisely because they think that genes play no privileged role in development.² In the view of many, it is a mistake to think that the important role of the genome in evolution also gives it a special role in development (Lehrman [1970], p. 136; Oyama [1985]; Griffiths [2005]; Mameli and Bateson [2011])—an idea that is encouraged by talk of genetic information or genetic representation.

Furthermore, the putative consumer of genetic representations is not a neat single mechanism, but consists of all the messy, interactive and temporally

¹ The teleosemantic approach to genetic information was pioneered by Sterelny *et al.* ([1996]), Maynard Smith ([2000]), and Jablonka ([2002]).

² See Griffiths ([2001]); Sarkar ([2004]); Godfrey-Smith ([2000, 2006a]); Levy ([2010]). In respect of whole organism phenotypes, see Stegmann ([2009]).

extended processes of development. Nor is it intuitive to see the developing individual as reading information contained in its zygotic DNA. After all, it's not as if there is any openness about which genes that individual could read, that is, respond to. Even if, as a result of selection, a range of different genotypes can be said to represent different environments and outcomes, an individual organism only has access to its own genotype, and not to a variety of different representations in that range. Godfrey-Smith argues that it is an 'adventurous' step to treat the complex processes of development that take zygotic DNA to whole-organism phenotypes as a single representation consumer:

Do we have any independent motivation to believe in these things [i.e. such consumers] at all, other than their enabling us to say something extra and intuitive about the content of the genetic message? (Godfrey-Smith [2007a])

Yet in some respects the way an organism reads genetic representations is on a par with the way it relies on adaptively relevant information detected in its environment. This paper argues that this symmetry furnishes a strong reason to treat some aspects of development as reading information carried by the genome, information that has been built up in phylogenetic time through the process of natural selection.

The account of genetic representation I developed in Shea ([2007a]) was entirely independent of how the idiosyncrasies of development unfold. It treated development as a black box supporting correlations between genotypes and phenotypes (in the range of environments in which selection occurred, against the background of the genetic variation then present). In that article, I also went further and conceded the Lehrman-inspired point that genetic representation not only does not depend on, but also tells us nothing about, how the processes of individual development are likely to unfold. This article re-examines that concession in the light of subsequent developments and concludes that, where development involves reading genetic representation (in the sense set out below), various defeasible inferences can be made about properties the developmental process is likely to have.

The article goes on to catalogue some further explanatory payoffs that derive from recognizing the existence of genetic representation of the kind identified by Shea ([2007a]): both in the genome (Section 5) and in development (Section 6). That serves to answer the kind of sceptic who admits that a coherent notion of genetic representation can be constructed, but doubts that it will do any useful explanatory work. These applications support the idea that genetic representation is a real phenomenon, and that identifying it is explanatorily useful. This is not a direct argument in support of the infotel semantic account of genetic representation of Shea ([2007a]), since the

explanatory utility of genetic representation is consistent with other accounts of its nature. However, the infotel framework does derive some indirect support from its being an account of genetic representation that can underpin these explanatory roles.

The theoretical framework behind the existence of genetic representation found in Shea ([2007a]) is summarized in Section 2. The wider term ‘inherited representation’ covers systems other than genes, if there are any, that have been designed to transmit phenotypes down the generations. Epigenetic inheritance and some forms of cultural inheritance are candidates. Section 3 argues that the way development reads genetic information is on a par with cases where development is designed to detect adaptively significant correlations in the environment. Section 4 argues that there are cases in which natural selection has been driven by the relative quality of environmental and genetic sources of information. That it is the relative quality of two sources of information that is driving evolution in these cases only comes into focus when we take the representational perspective. It follows that we have good reasons to recognize that genetic representations *are* read in the course of individual development.

2 Inherited Representations

This section summarizes the theory of representation in inheritance systems developed in Shea ([2007a]). I outline a representational framework, infotel semantics, then apply it to inheritance systems.

At least in simple systems, representational contents are used to explain behaviour. The ‘basic representationalist model’ (Godfrey-Smith [2006b]) applies when the system generating such behaviour divides into components: (i) a range of states that are sensitive to the environment (these will be turn out to be the *representations*); and (ii) a *consumer* system which acts on these tokens, producing a different type of behaviour in response to each type on which it acts. Often there is also a *producer* system, whose job is to generate the tokens on which the consumer acts (Figure 1).

Figure 1. The basic representationalist model.

There are various ways in which the tokens on which the consumer acts could have their content fixed. In some simple systems, appropriate contents are delivered by infotel semantics (Shea [2007b]), a modified version of teleosemantics (Millikan [1984]; Papineau [1987]). According to infotel semantics, content is fixed as follows:

Tokens of type R have content C if

- (a) there is a consumer system which is caused by a range of tokens, including tokens of type R , to produce a range of outputs, with a specific evolutionary function for each type of output, and where every token satisfies (b) to (d) with respect to some content;
- (b) R s carry the correlational information that condition C obtains;
- (c) an evolutionary explanation of the current existence of the representing system adverts to R s having carried the correlational information that condition C obtains; and
- (d) C is the evolutionary success condition,³ specific to R s, of the output of the consumer system prompted by R s.

Correlational information is a matter of types going with types. One object or system can have a range of properties, F_i , and which property it in fact has is informative about which property, G_i , some other system is likely to have.

Correlational Information

Object a 's having property F_1 carries correlational information about object b 's having property G_1 iff $P(G_1(b) | F_1(a)) > P(G_1(b))$

For infotel semantics, content is a matter both of how a representation is produced, and of how it is consumed so as to lead to behaviour. Infotel semantics naturalizes the intuitive idea that the consumer system treats R as a proxy for some fact about the world. When it receives R as input it behaves in a way that would be appropriate if C obtains.

The idea of a behaviour being appropriate in certain conditions is given in evolutionary terms. To use a standard example, the honeybee's nectar dance causes observer worker bees to fly off a certain distance in a certain direction to search for nectar. Consumer bees, observing a dance D of three waggles vertically have a disposition to fly off 250 metres in the direction of the sun, say. That behaviour has many evolutionary success conditions—conditions which enter into an explanation of how behaviour of that type in the evolutionary past led to survival and reproduction. Some conditions are very general (e.g. that the behaviour should promote fitness), but there is one evolutionary success condition that is specific to dances of type D : that there should be nectar 250 metres away in the direction of the sun. That

³ That is, the proximal evolutionary explanation of the survival and reproduction of the representing system adverts to C 's obtaining when R s were tokened.

Figure 2. The producer and consumers of a genetic representation.

delivers a Condition C in clause (d) above, dealing with the output side. At the input side, incoming bees are disposed to produce dances that do correlate with the location of nectar. When an incoming bee produces a dance of type D, of three waggles in the vertical direction, there is a good chance of nectar 250 metres away in the direction of the sun (Condition (b)). Furthermore, to explain how the whole system evolved, we need to advert to that correlation (Condition (c)). So the dance represents that there is nectar 250 metres away in the direction of the sun.

Infotel semantics vindicates the claim that zygotic DNA carries representational content as it is transmitted down the generations (Shea [2007a]). Figure 2 above shows the components corresponding to producer and consumers. Unintuitively, the producer is not a single organism, but a whole, temporally-extended episode of natural selection. Once selected, representations are then transmitted vertically down the generations, with multiple representation consumers, each constituted by the developmental systems of an individual organism. An analogy is the way cable TV gets transmitted down a street, with multiple branches coming off into individual homes, each with its own decoder. When they act on a selected genotype, transcription and translation are both part of the process of consumption, as are the complex interactions in which proteins fold, build structures, and carry out metabolic functions. The consumer is the long complex process

taking a selected genotype—through interaction with other parts of the organism, and with the environment, via the expression of multiple genes—to a phenotypic trait for which it was selected. By contrast, DNA replication is the basis of representation transmission (from one collection of zygotic DNA to its descendants).

The infotel framework applies as follows to zygotic DNA. A variety of different genotypes are found over evolutionary time (Condition (a)). Where a genotype G gives rise to a heritable phenotypic difference P , it may be selected. A phenotype P will be selected because of the way it interacts with some feature E of the environment (including existing features of conspecifics and of the organism itself).⁴ As a result, G acquires the evolutionary function of giving rise to P . There is an evolutionary success condition specific to performance of that function, namely that environmental factor E should still obtain (Condition (d)). Selection also has the effect of increasing the frequency of G in the population. Since that process occurs because E obtains, the result will be a correlation between G and E (Condition (b)), which explains the current existence of the whole system (Condition (c)). The upshot is that the zygotic DNA transmitted down the generations has semantic content. It indicates that environmental conditions are conducive to the phenotype, production of which caused that genotype to be selected. Infotel semantics will also attribute imperative contents to zygotic DNA in the same cases: the instruction to produce the corresponding phenotype (for simplicity, imperative content is not covered by Conditions (a) to (d) above).

The framework does not presuppose or require genetic determinism. It is consistent with the rich interactive complexity of the genotype-phenotype map. Development of the selected phenotype may require the presence of all kinds of contingent features of the environment. The framework only makes the uncontroversial assumption that there was a genotype-phenotype correlation at the time of selection. That correlation may subsequently be overlaid as products of the gene acquire a role in forming further phenotypic features, leading to pleiotropy, polygeny, and epistasis.

The genotype-phenotype correlation on which natural selection acts may be very long-distance. A gene may be selected for a phenotype that is not expressed until adulthood. It may also be selected for its effects in the world (beaver dams, bower-birds' bowers, etc.). Such extended phenotypes (Dawkins [1982]) can figure in the representational contents carried by genes. Genetic representation is not a matter of tight causal connection, like that seen

⁴ That is an idealization. There are doubtless many cases where a single phenotypic effect drives the selection of a gene. But a gene may also be selected for more than one of its phenotypic effects, each of which is selected in virtue of the way it interacts with one or more specific features of the environment. In such cases, the representational content carried by the gene will be correspondingly more complex.

between genes and their immediate protein products, but a matter of what a gene was selected for. Natural selection does not care how long and contingent the causal route between gene and phenotype is, as long as the phenotype was heritable in the range of environments in which it was selected.

In asexually reproducing organisms, the carriers of content are entire genomes, typed by genotype, not the just particular genetic polymorphisms between which individual episodes of selection occur. In sexually reproducing populations, it is appropriate to ascribe contents to individual genes.

Non-coding DNA will have content if it has selective functions, but not all the DNA transmitted down the generations has semantic content, e.g. neutral alleles or genes generated by random drift. One consequence of divorcing genetic representation from causal genetic determination is that genes that look to be strong causal determinants of phenotypic outcomes do not automatically count as representational (if, for example, they increase in frequency merely due to genetic hitchhiking). Shea ([2007a]) argues that such consequences, although at odds with some pre-theoretical uses of ‘genetic coding’, are virtues of the account.

Crucially, this framework only applies if DNA really does have the evolutionary function of transmitting phenotypes down the generations. That substantial but plausible empirical commitment is argued for elsewhere (Shea [2007a]). Bergstrom and Rosvall ([2010]) think it is obvious from its structure that DNA has been adapted for transmitting information. However, the DNA-based inheritance system may instead have arisen by chance, just happening to be particularly good at being an inheritance system. Maynard Smith and Szathmáry ([1995]) point to lines of evidence that it was the result of selection for its superior information-transmission qualities, but our framework does not depend upon these necessarily uncertain inferences about the origin of DNA, since there is a stronger source of evidence about function, in the way the mechanism has been improved by selection.

First, mechanisms of DNA proofreading and repair look like adaptations to improve transmission fidelity (Alberts *et al.* [2004], pp. 169–91). They are found in single-celled organisms, so cannot be explained away as adaptations for somatic cell inheritance. Second, it seems that the details of the way the triplet-amino acid code has been set up have been optimized to reduce the impact of common replication errors (as well as minimizing the effects of translation errors) (Haig and Hurst [1991]; Freeland and Hurst [1998]). Even if the origin of DNA was fortuitous, if these features are evidence of selection to improve its capacity for transmitting phenotypes, then DNA will thereby have acquired that function.⁵ If so, the mechanisms of development

⁵ That is, at the very least DNA has been exapted to play a central role in a mechanism that has the evolutionary function of transmitting phenotypes down the generations.

have the function of producing phenotypes in response to zygotic DNA, and so count as a *bona fide* consumer system. If there are other systems of inheritance with that metafunction, they too will carry inherited representations down the generations (for example, imitation in humans: Shea [2009], or epigenetic mechanisms: Shea *et al.* [2011], Shea [forthcoming a]). Having such a metafunction is a demanding constraint on the existence of inherited representations (Shea [2011b]).

3 Reading Genetic Representations

If genetic representation is the answer, what's the question? We can distinguish two broad questions that can be asked about an individual episode of development: why did it arrive at a particular outcome; and how did the process unfold? This section focuses on the former, arguing that genetic representation explains some of the cases in which the outcome matches a feature of the organism's environment. We return in Section 6 to questions about how the developmental process itself unfolds.

Sometimes a match between organism and environment is not adaptive and can simply be explained by causal processes occurring during the course of development. For example, brown bears on the mountainsides of Romania often have a blueish-purple muzzle in late summer. The shade of colouration rather closely matches the juice of the blueberries found on the same mountainsides. Since the match is not adaptive, once we observe that the juice causes the colouration, there is nothing more to be explained. (And sometimes an apparent match is just fortuitous so cannot be explained at all, say if the height in millimetres of a particular bear were to be the same as the height in metres of the mountain on which he lives.)

Our focus will be cases where the match between organism and environment is adaptive. That, too, can be achieved by causal interaction between the developing organism and its environment. The organism may have learnt a behaviour that is appropriate to its environment. Or it may have some other mechanism of adaptive plasticity which detects an environmental cue and produces an appropriate phenotype in response. For example, the water flea *Daphnia pulex* grows defensive armour if it detects a chemical correlate of predators in the water in which it develops.

In other instances, in contrast, the organism does not achieve the match by detecting environmental information. A striking example is provided by the ostrich, which is born with calluses in two places where its skin will touch the ground. Calluses are normally formed by mechanisms that cause skin cells to proliferate in response to friction. So the adaptive match between callus location and environment is explained by an environmental cue (friction at that location). In the ostrich, the developing embryo does not have access to any

environmental cue that would tell it where calluses would be adaptive. Instead, it appears that a gene or genes have been selected that set off the development of callosities in the embryo, triggering the process that is normally caused by friction (Waddington [1942]). What is striking about this case is that the mechanisms of development somehow lead to an outcome which adaptively matches the organism's (future) environment. Our explanandum is: of all the trajectories that development could have followed, how did it manage to hit on a trajectory that matches the individual's environment in this respect? How does the developing organism manage to narrow its uncertainty about the environment in which it will find itself? (Shorthand: how did the organism know the skin would be abraded just in those locations?—in a very thin sense of 'know'.)⁶

If Waddington's explanation is right, at some point a genotype G_C arose that produced the neonatal calluses, so that the difference between pre-callused and uncalled phenotypes correlated with a difference between genotypes (against the genetic background of that population, in the range of environmental conditions in which they were found). The pre-callused genotype G_C conferred a survival advantage and was selected.

According to the account of representational content outlined above, the genotype G_C will carry semantic content. As a result of selection, G_C carries correlational information about the content C : that there will be abrasion at a particular ventral location xyz :

$$P(C, |\text{an individual's genotype is } G_C) > P(C)$$

That correlation figures in the evolutionary explanation (Condition (c) above). On the output side, the mechanisms of development react to G_C by producing calluses at location xyz . There is an evolutionary success condition specific to that output, namely that the ostrich's skin is abraded at xyz (Condition (d)). It follows that G_C has indicative content: *there will be abrasion at ventral location xyz*. Infotel semantics also implies that G_C has imperative content: *produce calluses at ventral location xyz*. The indicative content can be false, say when the environment or the bird's anatomy changes and that part of the skin is no longer abraded, if development were then to carry on producing the calluses in response to G_C . The imperative content can also go unsatisfied if things go wrong during development and the organism bearing G_C fails to produce the pre-formed calluses (e.g. because of disease or malnutrition). In each case, the consumer of the representational content carried by the genome consists of the mechanisms of development. With that in hand, we can return to our explanandum: how does development narrow its antecedent uncertainty about the kind of environment the organism

⁶ Griffiths ([2008]) uses this formulation.

will find itself in? The answer is that, for the location of calluses in the neonatal ostrich, development relies on information carried by the genome.

Where the mechanisms of development have been designed to react to tokens because they carry correlational information about some adaptively relevant matter of fact, we can say the consumer is ‘reading’ information carried by those tokens without supposing that development is interpreting or understanding those tokens. Causally, the process of reading is just a matter of reacting to a token with an adaptively relevant response. Where there is a mechanism of adaptive plasticity designed to detect an environmental parameter in the course of development and rely on this parameter to produce an appropriate outcome, there is little resistance to the idea that development is narrowing its uncertainty by reading information in the environment. That is the standard story about normal calluses that result from abrasion. In such cases, if we ask how the organism ‘knows’ the relevant fact about the environment—the fact which it matches adaptively—the answer is that it detects the relevant fact in development. When selection is responsible for the information, there is much more resistance to the idea that information is being read in development. The later Lorenz ([1965]) argued for this position, working with a rather different concept of genetic information (Griffiths [2008]), but it never received widespread acceptance, despite the continuing pull of the question about adaptive match (Dawkins [1995], pp. 55–69).

One worry is that the distance between representation and outcome is too great. Many phenotypes emerge years after an organism begins developing. Nevertheless, the existence of long-range correlations between zygotic DNA and phenotypes, *at the time of selection*, is a staple of evolutionary theory. A second gap is the distance between the knowing thing and the thing known: the organism is picking up on correlational information generated through a process of selection that may have occurred in the distant past of phylogenetic time. That intuition is misplaced. Such correlations with the distant past are no more problematic than those relied on by geologists or palaeographers. A final worry is that the information generated by natural selection is only ever ‘known’ by a whole population or species (intuitively: by Mother Nature), not by individual organisms. The answer is that, as well as population-level information, natural selection generates correlations at the level of individual genotypes (more on which below).

In fact, there is just as much reason to see development as reading information in the genetic case. Rather than reacting to an environmental correlate of the adaptively relevant property so as to produce an appropriate phenotype, the developing organism reacts to a genetic correlate of the adaptively relevant property so as to produce an appropriate phenotype. In the genetic case, development ‘knows’ the relevant fact about the environment by reading the information carried by its genome. Both with the environmental cue and

with the selected genotype, there is a mechanism designed to react to correlational information and output an appropriate phenotype in response: a mechanism of adaptive plasticity in the first case, and the mechanisms of development taking DNA to phenotypes in the second case.

We need to distinguish development being designed to pick up on correlations in the environment from development depending causally on features of the environment (which is ubiquitous—light, nutrition, gravity, etc.). Development is only reading environmental information when it has been selected to produce different phenotypes in response to different environments. Normal skin development in the ostrich doubtless depends causally on many features of the environment. Suppose it needs a particular nutrient in the diet. Although ingestion by the mother of that nutrient would then correlate with production of the pre-callused phenotype, the nutrient is not being read as an environmental cue because presence of the nutrient was a fixed feature of all the normal environments in which the phenotype was selected (we can suppose).

Developmentally plastic outcomes will come to appear developmentally entrenched if the environment changes so that variation present at the time of selection is removed. Suppose predators of *Daphnia* were to proliferate so that their chemical traces were ubiquitous and the water fleas always grew defensive armour.⁷ Armour would then look just like a genetically selected phenotype that had gone to fixation. In terms of the proximal mechanisms of development, there may be no discernible difference between a mechanism of adaptive plasticity which now always encounters the same environmental cue on the one hand, and a genetic representation of the same matter of fact on the other. However, there is an important difference in the selective history of the two cases. With a gene gone to fixation, the source of the adaptive match between phenotype and environment has just been a process of selection. With a developmentally-entrenched environmental effect the story is more complex, with selection having generated information about an appropriate mechanism of adaptive plasticity, and environmental cues being responsible for the match between phenotype and environment in particular cases.

3.1 Do genes carry correlational information?

As promised, we return to the claim that, as a result of selection, genes carry correlational information. Consider a thought experiment. Imagine a temperature sensor in an organism that works as follows. A cell synthesizes two

⁷ Leave aside that there may then be selection pressure for genetic assimilation of the outcome—see below.

compounds in roughly equal quantities and secretes them into the cell membrane. Compound *A* is more stable below 20°C, say, and compound *B* is more stable above 20°C. After a while, a mechanism in the organism samples from the membrane and uses the result to determine a plastic phenotype (a temperature specific behaviour, morphological variant, developmental pathway, or some such). Since the ratio of molecules of *A* to *B* in the membrane will reflect the temperature, there are two ways the organism could use the compounds to determine its subsequent phenotype. It could sample many molecules and measure the ratio of *A* to *B*. Or it could just sample a single molecule, and produce the cool-appropriate phenotype if that molecule happened to be an *A*, and the warm-appropriate phenotype if that molecule were a *B*.

The first sampling method reflects the fact the population of molecules in the membrane carries information: the ratio of *A* to *B* in the population being high correlates with the environment's being cool. The second sampling method reflects the fact that the identity of a single molecule picked at random from that population also carries information. A single sampled molecule's being *A* correlates with the environment being cool. Correlation is stronger with a larger sample, but there still *is* a correlation at the level of the single molecule. Both correlations turn on conditional probabilities: what the temperature is likely to be, given various ratios of *A* to *B*; what the temperature is likely to be, given that the sampled molecule is an *A*. Our hypothetical biological thermometer could work either way.

Natural selection on genes gives rise to a big, temporally-extended detection mechanism, one which works in just the same way as the biological thermometer in our thought experiment. Random mutation produces a population in which there are two heritable phenotypes, P_A and P_B (based on the transmission of genes *A* and *B*). Suppose environmental temperatures below 20°C are conducive to P_A , and those above 20°C to P_B . If other factors remain stable, natural selection will lead to a population in which the ratio of gene *A* to gene *B* carries correlational information about the environmental temperature. One gene may even have gone to fixation. As an experimenter, if we know the heritability facts, we can read off from the population-level statistics facts about the likely temperature of the environment in which selection has occurred. But notice, as with our hypothetical biological thermometer, there are also correlations at the level of individual genomes. If we pick a genome at random, the fact that it carries gene *A* is also an indicator that the selective environment was below 20°C. Indeed, if the selection is strong enough to go to fixation, then the correlational information carried by a single genome is just as reliable as that obtainable from a larger sample. The central dogma is so important because it ensures that information built up in the genome by this process is not degraded by things that occur in individual ontogeny.

So genes and genotypes carry correlational information in just the same sense as environmental cues carry correlational information. The only difference—and it is an important difference—is that the information carried by a gene arises from a process that takes place over phylogenetic time, whereas the information in the environmental correlate subsists over much shorter timescales. Inherited representations correctly indicate how things are now only if the environment has not changed in the relevant respect for a very long time (Category *A*). Setting a developmental trajectory based on an environmental cue will only be useful if the environment does not fluctuate in the relevant respect at the ontogenetic timescale (Category *B*). Other cues read by an organism are extremely transient, for example a cue that there is a fly passing by right now (Category *C*). In each case, the timescale of the relevant information is matched by features of the mechanism which collects it. Information that is stable over phylogenetic timescales is collected by a process of natural selection. Information that is stable over ontogenetic timescales is made use of by developmental programs to drive phenotypic plasticity. And transient, non-portable information is usually exploited by physiological mechanisms that can react at a timescale of seconds or milliseconds.

4 Selection Between Genetic and Environmental Information

4.1 Modelling

So far, I have argued that we should recognize that there are genetic representations, of the kind specified above, and that they are read in individual development. What does this perspective add? When information theory has been brought to bear in biology, sometimes it merely offers a new way of talking about things biologists understand well already (e.g. the effect of natural selection on evolutionary dynamics can be re-expressed as maximizing Fisher information: Frank [2008]). Other times it provides a better way of understanding known phenomena (e.g. the fact that information is not subject to intrinsic conservation laws throws new light on the cooperation involved in some higher levels of organization: Lachmann *et al.* [2000]). Here we are concerned, not with formal information theory, but with the merits of a representational treatment of inheritance systems that captures the semantic information that goes with selective functions. Our claim will be that this framework is not only compatible with the empirical facts but, like Lachmann *et al.* ([2000]), sheds new light on some biological processes.

I do not claim that the representational approach is indispensable. Since our account of genetic representation is a naturalistic reduction, explanations in terms of genetic representation can always be dispensed with in favour of the properties to which it reduces. Reduction of a property is not elimination, but

rather vindication of its existence, so the critical question is whether the reduced property has any important explanatory work to do. I aim to show that appeals to genetic representation offer improved explanatory purchase on some biological phenomena. In this section, I argue that the known phenomena of environmental and genetic phenotype determination are better explained when we recognize the role of genetic representation. The symmetry between reading environmental cues and reading genetic representations is central to understanding such cases. So the recognition that genetic representations are read in individual development has an important explanatory payoff.

Our examples concern adaptive polymorphisms. For instance, some species have two morphological variants in two different spatial environments, e.g. forming a bush on high, cool mountainsides but growing as a tree in warm valleys. The difference may be due to a genetic difference or to adaptive phenotypic plasticity, with the developing seedling directly detecting a cue of its likely environment (e.g. temperature). Similarly, colour variation is genetically caused in many cases (Gray and McKinnon [2007]), but may be environmentally induced (Gotthard *et al.* [2009]). Other polymorphisms concern life history strategy (e.g. when to flower, how quickly to develop) and behavioural dispositions (e.g. on which host to lay eggs). In aphids, there are examples of both genetic and environmental determination of the choice between winged and wingless forms. There are parthenogenetic females in which development switches from a wingless to a winged form based on an environmental cue that the quality of the host plant is declining. In contrast, there are dimorphic male aphids in which the difference between winged and wingless forms depends upon a genetic difference (Braendle *et al.* [2006]).

Environmental and genetic phenotype determination are standardly taken to be different phenomena. It has been noted that the accuracy of available information about the environment affects whether or not a polymorphism is environmentally cued (West-Eberhard [2003]; Hazel *et al.* [2004]). Likelihood of genetic morph determination depends on factors like the extent of genetic isolation of two subpopulations (Sultan and Spencer [2002]). Such factors were not seen in informational terms until recently. Leimar *et al.* ([2006]) investigated the circumstances in which the choice between genetic and environmental morph determination can itself be adaptive. They modelled an idealized species in which there are two available morphs, set by a developmental switch, where the species can evolve towards setting that switch in reliance on an environmental cue, a genetic cue, or a combination of the two (Figure 3).

Leimar *et al.* find, unsurprisingly, that the environmental cue is given more weight when it is more reliable. Their innovation was to show that reliability of the genetic cue drives development towards relying on the genotype for morph determination. One factor that makes the genetic information more reliable

Figure 3. A developmental switch reliant on genetic and/or environmental cues (adapted from Leimar *et al.* [2006]).

is the existence of strong selection pressure. Another is restricted gene flow between patches in the environment. Others have recognized the relevance of these factors, but without observing their effect on the quality of available genetic information. To the extent that subpopulations are isolated genetically, the informational signal generated by selection is not diluted by genes arriving from environments where the relevant environmental parameter is different. Strong selection pressure generates correlational information more quickly. Another factor we might add is population size. In a sexually reproducing population, random drift will degrade genetic information more rapidly in a small population. Leimar ([2009]) locates the modelling work in a wider theoretical framework which explicitly recognizes that genetic information is read in the course of individual development:

[...] genetic cues in the form of allelic variation at polymorphic loci can play similar roles as environmental cues in providing information to the developmental system about coming selective conditions. (Leimar [2009], p. 125)

Interestingly, in some conditions the models in (Leimar *et al.* [2006]) evolve to a mixed solution in which both genetic and environmental cues play a role in setting the developmental switch. Such mixed solutions are not just a staging post on the evolutionary trajectory towards either genetic morph determination (thus forming two genetically distinct subpopulations) or on environmental morph determination (the mechanism of adaptive plasticity then being an adaptation that is shared by both subpopulations). The modelling shows that the mixed solution can itself be a final stage—an adaptation which optimizes reliance on the genetic and environmental cues that are available.

This modelling work strengthens our argument that the reasons for seeing environmental morph determination in informational terms carry over to

genetic morph determination. But it goes further, showing that evolutionary processes can be driven by the relative quality of these two sources of information. This is a process that can only be understood in terms of the information carried by the two kinds of cues. Describing either the environmental or the genetic cause of morph determination in non-informational terms would lose sight of the fact that an adaptive choice has been made based on the relative quality of two different sources of information: the environmental cue that covaries in real time with climatic conditions, and the genetic cue which carries correlational information in virtue of selection over many generations. That answers sceptics by showing the important explanatory work that can be done by acknowledging that there is genetic representation and recognizing that it is read by development.

Do the Leimar *et al.* models give a reason to discern information at the level of the individual organism? Jablonka and Lamb [2005] compare information channels in terms of their properties for a population. Similarly, the selection modelled by Leimar *et al.* occurs at the level of the population. Couldn't all the points about information be made just at the level of the population as a whole? No, because reacting to the environmental cue is a matter of how individuals read information in the course of their individual development. We could see selection as choosing between an obligate developmental trajectory and an adaptively plastic one, where only the latter involves information being read by an individual in development. However, that would lose sight of the fact that selection is based on the quality of two sources of information. With environmental cues, we have information quality for an individual. The most direct way of seeing selection between that and genetic information also requires us to quantify the quality of the other (genetic) source of information for the individual. Such symmetry of treatment becomes even more important when the system evolves to the mixed solutions discussed in Leimar *et al.* ([2006]).

4.2 Empirical applications

Leimar's modelling results are vindicated by some real empirical examples. The point of these examples is that selection based on the quality of the information in environmental and genetic representations, illustrated in the models, has really occurred in practice.

Lind and Johansson ([2007]) studied variations in developmental time in the common frog, *Rana temporaria*, sampled from fourteen different islands. Frogs should develop more quickly in areas where ponds dry up more quickly. They found two effects. First, across the fourteen islands, mean developmental time on an island correlated with the mean drying out time of ponds on the island. Second, frogs from islands with greater pond variability showed

greater plasticity in the adaptation of their developmental time to simulated pool drying in the laboratory. Lind and Johansson explained the existence of these two effects by reference to the modelling in Leimar *et al.* ([2006]): genetic cues are more reliable on homogeneous islands, leading to decreased plasticity and stronger reliance on the genetic cue of the mean drying time for the island, whereas environmental cues are more important on islands where ponds vary more in their speed of drying.

Other results that have been explained in the same way include geographical patterns in butterfly growth rates (Nygren *et al.* [2008]) and butterflies' choice of host plant on which to lay their eggs (Janz *et al.* [2009]). An interesting case of environmentally-cued polymorphism is the colour dimorphism of larvae of the butterfly *Pararge xiphioides* (Gotthard *et al.* [2009]). It seems that there is a greater advantage to being dark coloured in high density populations, partly for reasons of camouflage. Gotthard *et al.* ([2009]) explain the observed phenotypic plasticity by the fact that, given the ecology of the species, the most reliable cues of population density are environmental.

In principle, there could be genetic cues that act in a similar way, given an appropriate population structure. For example, when there is selection for a phenotypic variant that is adaptive only at low frequency in a population, selection will generate the correlational information that the frequency of that phenotype in the population is low. On our account it follows that the genome represents that that particular phenotype is found only at low frequency in the population. This example illustrates that genetic representation need not be only restricted to facts about the external environment, but may also concern properties of conspecifics, and of the population in which the individual is likely to develop.

Leimar ([2009]) discusses further empirical observations that he argues are best explained in terms of information in environmental and genetic cues. For example, he argues that the parthenogenetic female aphids, in which wing dimorphism is environmentally determined, have access to reliable cues of plant quality, whereas in aphids where males are only produced for short periods in the autumn, no reliable cue of host quality is available to the male, explaining reliance on a genetic cue for wing dimorphism in such cases (p. 132).

Other examples can be drawn from the literature on genetic assimilation and genetic accommodation. In genetic assimilation, an environmentally induced phenotype becomes genetically fixed. The ostrich calluses discussed above are an example, on the reasonable assumption that they were once environmentally caused (Waddington [1942]). It is becoming clear that environmentally induced variation is a significant source of the novel phenotypic variation on which natural selection acts (Gilbert and Epel [2009], pp. 369–402). Genetic assimilation is selection of such environmentally

induced variations, often without genetic mutation but relying instead on pre-existing (often cryptic) genetic variants in the population which are more canalized towards the adaptive outcome. Genetic accommodation is evolution in the opposite direction, towards phenotypic plasticity. When a new trait is induced by an aspect of the environment, but is adaptive only in some environments, there can be selection pressure for plasticity—for producing the trait only in those environments. For example, artificial selection in the hornworm *Manduca sexta* led to genetic accommodation of a plastic phenotype (Suzuki and Nijhout [2006]).

When will there be genetic assimilation and when genetic accommodation? The representational framework offers a perspective for unifying these phenomena. Take some developmental process that is causally sensitive to a variable feature of the environment. If that variation correlates with an adaptively relevant phenotypic outcome, then there will be selection pressure to use that cue to determine the outcome (genetic accommodation). If not, there will be selection pressure to insulate development against the variation (genetic assimilation). For example, the rate of gonad growth is temperature-dependent, making sex determination in cold-blooded animals potentially dependent on the ambient temperature. If temperature is a sign of something adaptively relevant, for example the sex ratio that is most adaptive at that time of year, then there will be selection pressure for genes which channel that variation towards the adaptive outcome (Charnov and Bull [1977]; Warner *et al.* [2009]; Pen *et al.* [2010]).

The representational framework has the merit of underpinning modelling and explanation of these evolutionary transitions between environmental and genetic cues, since it treats both as instances of the same phenomenon: acting on a source of information for its fitness value. The framework does not yet rely on mathematical information theory. However, a recent paper shows that there are strong connections between the fitness value of information in uncertain environments, as discussed here, and the mutual information carried by a developmental cue, as measured by information theory (Donaldson-Matasci *et al.* [2010]).

To summarize so far: the combination of modelling work with its exemplification in real systems strongly supports the claim that genetic representations are read in individual development.

4.3 Maternal effects

The representational framework also applies to cross-generational phenotypic plasticity (Sultan [2000]), for example as mediated by maternal effects (cf. Badyaev and Uller [2009]; Shea [forthcoming a]). In such cases, the organism does not detect the environmental parameter itself, but relies on its

mother to do so, making use of a cue sent by the mother that correlates with facts about the mother's environment (the cue is often communicated by chromatin marking or some other epigenetic mechanism). Galloway and Etterson ([2007]) studied maternal effects in the herb *Campanulastrum americanum*. They found that maternal light environment (light gap versus understory) influences which life history strategy offspring adopt (annual versus biennial). This transgenerational plasticity is adaptive when the offspring grow in the same light environment as their mother: being annual is the best strategy in a light gap, and biennial in the understory.

Why is this effect mediated by maternal effects? We can answer that question in terms of the reliability of the different sources of information potentially available to inform development (Uller [2008]). Presumably, the seed is not in a good position to detect directly whether it will be in a light gap or understory, hence whether to germinate in the spring or the autumn (biennial or annual, respectively). On the other hand, genetic information will not carry a good signal either, since pollen will spread from a much larger area than the offspring's immediate environment, so will come indiscriminately from fathers in both understory and light gaps. However, with limited seed dispersal, the offspring's environment is likely to match that of its mother. The time course of fluctuations between understory and light gap is probably too fast to be registered by selection on maternal genes. But maternal effects can be a good source of information. Therefore, we can understand why life history strategy is determined by maternal effects in *C. americanum* in terms of the plant having evolved to act upon the most reliable source of information for choosing between polymorphic states. Reliance on transgenerational phenotypic plasticity is well-explained in terms of the reliability of the various sources of information available to development: genetic, epigenetic, and environmental (Uller and Pen [2011]).

Since the investigation has only just begun, these cases suggest that there will be a large number of empirical phenomena that are best explained in terms of the genetic, epigenetic, and environmental information that is read in the course of individual development.

5 Genetic Representation and the Genome

5.1 Information capacity of organisms' genomes

The hypothesis that there is genetic representation, with DNA being adapted to transmit such representations down the generations, can throw explanatory light on further features of the genome. In this section, I briefly mention four. Recall that I do not claim that these features can only be explained in representational terms. Rather, I aim to show that recognizing the existence of

genetic representation (as I have characterized it) throws new light on some biological phenomena. I start with the observation that the genomes of complex organisms have a very high information capacity.

The information capacity of a channel or store of information is a matter of the extent to which it can adopt a large variety of differentiable, readable states (often measured in terms of the number of binary values, or bits, which could be encoded). Why, exactly, should we expect the information capacity of the genome of a complex organism to be high? The answer is that a high information capacity is needed if a genome is to represent a long, conjunctive content; and developing into a complex organism with a large number of genetic adaptations requires a genetic representations with long, conjunctive contents. There is a tendency to move rather quickly here, so it is worth spelling out the reasoning with some care.

Two caveats: First, there is no simple correlation between the complexity of an organism and the length of its DNA. The proportion of non-coding DNA and the frequency of chromosome duplication vary enormously. The presence of introns in eukaryotes further complicates any direct relationship. But adaptive complexity imposes a minimum, and that minimum already requires genomes to store a very large amount of information, even if there are many reasons why actual genomes are even larger. Second, we are not here asking why DNA is so good at storing information, as discussed in Section 2 above. Our question is why the genomes of complex organisms are so long. The answer is uncontroversial, but the point here is to show it is best understood in informational terms.

The picture is clearest with selection on individual genes. When a new gene *G* arises (perhaps by neutral duplication and then mutation), and is selected for some novel beneficial phenotypic effect *P*, the frequency of *G* will increase, perhaps going to fixation. As a result, as we have seen, *G* carries indicative content about the selective environment, and the imperative content that development should produce phenotype *P*. *G* then becomes part of the genetic background against which further adaptations can occur. To build up further information through new episodes of selection, without degrading the information carried by *G*, the new information must be accumulated at other loci. So if a genome is to carry a very large amount of information generated by very many episodes of selection, it must have a very high storage capacity. This is information capacity at the input side: the ability to store correlations produced by a long sequence of different episodes of selection.

Complexity on the output side—in the organism that results from development—also requires the genome to store a large amount of information.⁸

⁸ That complexity implies that the organism has a high informational capacity, which is a further explanandum (cf. Godfrey-Smith [1996]).

When new genes are selected against the background of genomes containing G , such genomes must continue to give rise to the phenotype P if continuing use is to be made of the information carried by G . Even if the expression of G is causally involved in generating subsequently selected phenotypes, selection for such phenotypes must take place on different genes if the process is not going to interfere with production of P . It follows that development of an organism displaying adaptive complexity—that is, a large number of adaptive phenotypic traits—also requires the genome to store a large amount of information.

The inference from adaptive complexity to DNA storage capacity is equally justified when selection occurs at a higher level than the individual gene, although here things are a little more tricky. In asexual organisms, selection is occurring between whole genomes. Natural selection increases the frequency of some genotype G^* compared to other genotypes G in the population, where being of genotype G^* is a property of the whole genome (even if the factor differentiating G^* from other genotypes G is a single locus or a single nucleotide). The same is true in sexual species when there is a major reorganization of the genome (e.g. the chromosome duplication common in plants). There again the types between which selection occurs are entire genomes.

In these cases, too, adaptive complexity requires a high storage capacity. If the information that has been garnered by selecting genome G^* against alternative genomes G is not to be lost, subsequent episodes of selection must act on subtypes of G^* : G^*_1 and G^*_2 say. If old genotype G were subsequently selected over G^* , that would destroy the information generated by the original episode of selection. The picture is of G^* being selected over existing genotypes for some phenotype P_1 , and then G^*_1 being selected over G^*_2 for some second adaptive phenotype P_2 , so that G^*_1 comes to carry the conjunctive imperative content *produce P_1 and produce P_2* , and a corresponding conjunctive indicative content that *the environment is conducive to P_1 and the environment is conducive to P_2* . To be cumulative, further selection must then take place on sub-subtypes of subtype G^*_1 , and so on. If the message is to be cumulative, the genome must support a very large number of differentiable types, without going back to types that have previously been selected against. That requires a very high information capacity.

Therefore, in both cases, adaptive complexity requires an organism's genome to have a high information capacity. Contraposing, we can explain the high information capacity of extant genomes by appealing to their role in transmitting down the generations the information needed to produce organisms with complex adaptive phenotypes.

5.2 Many amino acids, few nucleotides

A second question susceptible to an informational treatment is as follows. Why are there twenty amino acids but only four nucleotides? Maynard Smith and Szathmáry ([1995]) offer an explanation in terms of adaptations for making use of information. They argue that there was a time when RNA acted as a *ribozyme*: both as an enzyme for metabolism and as a basis of replication. In that world, the number of nucleotides deployed was the result of a cost-benefit trade off. Having more nucleotides would increase the specificity of the enzymatic activity of ribozymes, hence their metabolic efficiency, but at the cost of suffering more errors in replication, due to decreasing the specificity of base-pairing. Given the high cost of replication errors, which would degrade the information available to all future generations, this trade off settled on a low number of nucleotides, privileging replication fidelity over enzymatic specificity.

Maynard Smith and Szathmáry argue that the ribozyme world was a precursor to a system, still based on RNA, in which replication and enzymatic activity became separated. RNA coded for amino acids, which did all the enzymatic work, with RNA serving just as a basis of replication. The invention of coding for amino acids allowed greater enzymatic specificity to be achieved without compromising the specificity of the base-pairings used in replication. Errors in translation are more tolerable, since a non-functional protein can be discarded and the RNA translated again. Such errors do not degrade the quality of information available to future generations. Accordingly, once a distinction was established between information carrier and enzymatic agents the triplet code could make use of a far larger library of enzymatic agents (amino acids). The cost of having a larger library is that more errors are made in translation (the wrong amino acid is produced from a nucleotide triplet), but that cost is much lower than the cost of replication errors which constrained the same trade off in the ribozyme world. The result is that life makes use of many more amino acids than nucleotides.

Theorizing about the early origins of life is necessarily uncertain. Maynard Smith and Szathmáry's theory is useful for our purposes, not because of the strength of the evidence for its conclusion, but because it illustrates the power of the informational approach to throw light on important problems. The trade off they describe could be captured without mentioning information, as an adaptation for evolvability. However, their innovative hypothesis can be expressed in the most clear and satisfying way when we see it in terms of information: the evolution of the triplet code was the invention of a channel of information, i.e. a process of information transmission that is separate from enzymatic/metabolic activity.

5.3 A function of sex

A third example is the function of sex. In an asexual population, genetic information generated by the process of natural selection is transmitted within the lineage in which the selection occurred (leaving aside lateral gene transfer). Useful information about the selective environment gets confined within the branch where it was first discovered. In a sexually reproducing population, that information can spread throughout the population. One possible explanation of the adaptiveness of sexual reproduction is that it allows individuals to have access to useful genetic information drawn from a wider field. Where there is strong selection pressure (high quality information), correlational information accumulated in the genome through selection in a small subpopulation can quickly spread to the rest of the species.

That is on the plus side, but on the minus side is intra- and inter-genomic sexual conflict. A gene can be selected that has negative effects on one sex provided it has sufficiently positive effects on the other. This process degrades the quality of the information available to the disadvantaged sex.⁹ A potential solution is a sex-specific disposition to ignore such degraded information (say by switching off the relevant gene).

Sex may also be important in aiding DNA repair based on complementary copies of genes, thus improving the fidelity of vertical information transmission. The origin and function of sex is a highly contentious topic, but it is instructive to notice that an informational treatment may have a role to play.

All these cases are advanced tentatively, since we must be wary of excessively optimistic adaptationism. However, that is a criticism that applies to both informational and non-informational treatments of the phenomena above. The important point for our purposes is that, to the extent that one or more of these phenomena turn out to be adaptations, although the representational framework is not indispensable, recognizing the existence of genetic representation affords a particularly satisfying explanation.

6 Explaining Further Aspects of Development

6.1 Canalization against environmental variation

In the next two subsections, we see that the representational framework casts light on further aspects of development. The point is not that these phenomena could only be predicted from the informational perspective, but rather that they can be explained in informational terms, thus supporting our thesis that genetic representations really are read in development.

⁹ Thanks for Tobias Uller for this suggestion.

Figure 4. Genetic transmission mapped onto the encoder, channel, and decoder of information theory.

This subsection shows that an informational take on development should lead us to expect genetically based adaptations to become canalized against environmental variation. This treats development as a black box. Subsection 6.2 gets inside the black box and uses an informational perspective to throw light on structures that are found within the mechanisms of development. The standard framework of information theory (Cover and Thomas [2006]) maps onto genetic transmission as shown in Figure 4.

DNA replication down the generations constitutes a transmission channel, with each developmental system (branching off from it) being a decoder. The overall aim of the communications engineer is that the estimate of the message $\hat{W}$ should match the message W with high probability: $P(\hat{W} \neq W)$ should be low. So far, we have looked at factors that affect the noise introduced by the channel (e.g. the fidelity of genetic replication) and the quality of information produced by the encoder (e.g. strength of selection pressure, genetic isolation of subpopulations). A further desideratum is to minimize the amount of noise introduced by the decoder: the entropy between signal and decoded message, $H(\hat{W}|Y^n)$, should be low. The mutual information between genotype and phenotype is negatively related to this entropy term:

$$I(\hat{W}|Y^n) = H(\hat{W}) - H(\hat{W}|Y^n)$$

For a given distribution of phenotypes—fixing $H(\hat{W})$ —reducing $H(\hat{W}|Y^n)$ will increase $I(\hat{W}|Y^n)$: the mutual information between the genotypes that can go into the developmental process and the phenotypes which result (roughly the strength of the genotype-phenotype correlation). For there to have been natural selection on genes—i.e. for the encoder to have operated in the first place—a phenotype must have been heritable, so there must have been mutual information between genotype and phenotype at the outset. But these considerations go further.

When a genetically-selected phenotype has gone to fixation, there will no longer be any genetic variance that is relevant to the trait. The noise introduced by the decoder can still be reduced (i.e. the mutual information between

genotype and phenotype can be increased), by reducing the noise caused by variations in the environment, that is, by canalizing development of the phenotype against environmental variation. In short, the information engineer's perspective predicts that we should expect genetically selected traits that have gone to fixation to become further canalized against environmental variation (Shea [forthcoming b]). The invariance which is a good thing for fitness turns out to be equivalent to the communication engineer's desideratum of reducing the noise introduced by the message decoder. From the point of view of selection, canalization increases fitness by increasing the range of circumstances in which a useful phenotype is produced. From the point of view of information theory, canalization is an adaptation to reduce the noise introduced by the decoder.

The informational perspective on invariance and canalization is illuminating in its own right. It also allows us to get inside the black box of development—to explain aspects of the way the decoder is implemented (i.e. of the way ontogenetic processes actually do the job of turning zygotic DNA into whole organism phenotypes).

6.2 An informational function for the nuclear membrane?

In organisms that pass through developmental stages, some of the information in the genome is not used until those later stages. The stability and inertness of the DNA molecule keeps that information from being degraded during the course of development. In multicellular organisms, some of that information needs to be preserved when cells divide. The replication fidelity of DNA in germ line inheritance does double duty here, acting as an effective means of preserving information in somatic cell inheritance as well.

Consider for a moment what would happen if all DNA were converted into proteins in the zygote and then discarded, with somatic cell inheritance achieved by copying protein samples, and stages in development occurring by the regulation of protein activity (cf. Godfrey-Smith [2000]). That is doubtless a fanciful scenario. We can see immediately that it would be a much less reliable way to generate whole organism phenotypes. The separation between information transmission and metabolic activity that is so useful in germ line inheritance (discussed above) is also a very effective means of preserving information until it is needed in the course of individual development. In organisms with developmental programs, correlations between genotypes and phenotypes are made more robust because genetic information is preserved in DNA throughout the course of development (which is not to say that it is an adaptation for that purpose).

In eukaryotes, the fidelity with which that information is preserved is improved by a further innovation, the nuclear membrane. By separating the site

of transcription into RNA within the nucleus from the metabolic activity that occurs after translation in the extra-nuclear cytoplasm, genetic information can be more effectively preserved during development. Our suggestion is that the nuclear membrane may have evolved for an informational reason: to increase the strength of genotype–phenotype correlations in complex organisms. This hypothesis could be tested by comparing the mutation rates in somatic cell inheritance between prokaryotes that show some cell differentiation and very simple eukaryotic organisms with similar amounts of cell specialization.

There is a puzzle about why prokaryotes have never evolved complex multicellularity (Maynard Smith and Szathmáry [1995], pp. 223–4), given that they have mechanisms of gene regulation and have invented mechanisms that could be useful for cell heredity (e.g. host-induced modifications passed on through chromatin marking). Part of the explanation may be that the development of complex multicellularity faces an informational problem, that of preserving genetic information until it is needed in development. The nuclear membrane may have arisen as an adaptive solution to that problem, being selected because of the greater capacity for cell differentiation and developmental stages in organisms that have it.

The informational perspective generates further insights into the mechanisms of somatic cell inheritance, which there is no space to discuss here (Shea *et al.* [2011]). The nuclear membrane example must suffice as an illustration of how the informational perspective can help explain why the internal mechanisms of development—developmental programs, somatic cell inheritance, etc.—take the form that they do.

7 Conclusion

Genetic representation gives the genome a special status. Yet the more closely we examine the causal unfolding of ontogenesis, the clearer it is that genes play no special causal role. They have no special causal force—they in no way causally determine the outcome of development. Ontogenetically, genes seem to be on a par with all the other causally relevant factors, epigenetic, environmental, and cultural. As a result, even those who accept that there is genetic representation think that its explanatory utility must be confined to organizing facts about evolution and selection. Genetic representation cannot be relevant to explaining individual ontogeny.

This article makes the case for the opposite view: genetic representation can help explain facts about individual ontogeny. The heterodox view that genetic representations are read by the mechanisms of individual development is no more objectionable than the idea that mechanisms of adaptive plasticity read informational cues in the environment. Indeed, to understand the way

evolution has selected between individuals who rely on these two sources of information to different degrees—and there do seem to be examples in many species—we have to understand both genetic and environmental factors in informational terms. Doing so has further explanatory bite. It allows us to explain various facts about the genome. More surprisingly, it allows us to explain some of the processes that occur within individual development, like the way eukaryote cells sequester DNA away behind a nuclear membrane. These considerations amount to a powerful answer to scepticism about the explanatory utility of inherited representations.

Acknowledgements

The author would like to thank Carl Bergstrom, Peter Godfrey-Smith, Paul Griffiths, Maureen O'Malley, David Papineau, and an audience at the University of Exeter for helpful discussion of this material; and Cecelia Heyes, Russell Powell, Tobias Uller, and three referees for this journal for detailed comments on earlier versions of the article.

Funding

John Fell OUP Research Fund, Oxford Martin School, Somerville College (Mary Somerville Junior Research Fellowship) and Wellcome Trust (grant number 086041 to the Oxford Centre for Neuroethics). Funding to pay the Open Access publication charges was provided by the Wellcome Trust.

*Faculty of Philosophy and Somerville College
University of Oxford, 10 Merton Street, Oxford OX1 4JJ, UK
nicholas.shea @philosophy.ox.ac.uk*

References

- Alberts, B., Bray, D., Hopkin, K., Johnson, A., Lewis, J., Raff, M., Roberts, K. and Walter, P. [2004]: *Essential Cell Biology*, 2nd edn., New York & Abingdon: Garland Science.
- Badyaev, A. V. and Uller, T. [2009]: 'Parental Effects in Ecology and Evolution: Mechanisms, Processes and Implications', *Philosophical Transactions of the Royal Society B*, **364**, pp. 1169–77.
- Bergstrom, C. T. and Rosvall, M. [2010]: 'The Transmission Sense of Information', *Biology and Philosophy*, **26**, pp. 159–76.
- Braendle, C., Davis, G. K., Brisson, J. A. and Stern, D. L. [2006]: 'Wing Dimorphism in Aphids', *Heredity*, **97**, pp. 192–9.
- Charnov, E. L. and Bull, J. J. [1977]: 'When is Sex Environmentally Determined?', *Nature*, **266**, pp. 828–30.
- Cover, T. M. and Thomas, J. A. [2006]: *Elements of Information Theory* (2nd edn), Hoboken, NJ: John Wiley & Sons.

- Dawkins, R. [1982]: *The Extended Phenotype*, Oxford: Oxford University Press.
- Dawkins, M. [1995]: *Unravelling Animal Behaviour*, 2nd edn., Harlow: Longman.
- Donaldson-Matasci, M. C., Bergstrom, C. T. and Lachmann, M. [2010]: 'The Fitness Value of Information', *Oikos*, **119**, pp. 219–30.
- Frank, S. A. [2008]: 'Natural Selection Maximises Fisher Information', *Journal of Evolutionary Biology*, **22**, pp. 231–44.
- Freeland, S. J. and Hurst, L. D. [1998]: 'The Genetic Code is One in a Million', *Journal of Molecular Evolution*, **47**, pp. 238–48.
- Galloway, L. F. and Ettersson, J. R. [2007]: 'Transgenerational Plasticity is Adaptive in the Wild', *Science*, **318**, pp. 1134–6.
- Gilbert, S. F. and Epel, D. [2009]: *Ecological Developmental Biology*, Sunderland, MA: Sinauer Associates.
- Godfrey-Smith, P. [1996]: *Complexity and the Function of Mind in Nature*, Cambridge: Cambridge University Press.
- Godfrey-Smith, P. [2000]: 'On the Theoretical Role of "Genetic Coding"', *Philosophy of Science*, **67**, pp. 26–44.
- Godfrey-Smith, P. [2006a]: 'Information in Biology', in D. Hull and M. Ruse (eds), *The Cambridge Companion to the Philosophy of Biology*, Cambridge: Cambridge University Press.
- Godfrey-Smith, P. [2006b]: 'Mental Representation, Naturalism, and Teleosemantics' 2006, in D. Papineau and G. Macdonald (eds), *New Essays on Teleosemantics*, Oxford: Oxford University Press.
- Godfrey-Smith, P. [2007a]: 'Innateness and Genetic Information', in P. Carruthers, S. Laurence and S. Stich (eds), *The Innate Mind: Foundations and the Future*, Oxford/New York: Oxford University Press, pp. 55–68.
- Godfrey-Smith, P. [2007b]: 'Conditions for Evolution by Natural Selection', *Journal of Philosophy*, **54**, pp. 489–516.
- Godfrey-Smith, P. [2011]: 'Senders, Receivers, and Genetic Information: Comments on Bergstrom and Rosvall', *Biology and Philosophy*, **26**, pp. 177–81.
- Gotthard, K., Berger, D., Bergman, M. and Merilaita, S. [2009]: 'The Evolution of Alternative Morphs: Density-Dependent Determination of Larval Colour Dimorphism in a Butterfly', *Biological Journal of the Linnean Society*, **98**, pp. 256–66.
- Gray, S. M. and McKinnon, J. S. [2007]: 'Linking Color Polymorphism Maintenance and Speciation', *Trends Ecology and Evolution*, **22**, pp. 71–9.
- Griffiths, P. E. [2001]: 'Genetic Information: A Metaphor in Search of a Theory', *Philosophy of Science*, **68**, pp. 394–412.
- Griffiths, P. E. [2005]: 'The Fearless Vampire Conservator: Philip Kitcher, Genetic Determinism, and the Informational Gene', in E. M. Neumann-Held and C. Rehmann-Sutter (eds), *Genes in Development: Re-reading the Molecular Paradigm*, Durham, NC: Duke University Press.
- Griffiths, P. E. [2008]: 'The Distinction Between Innate and Acquired Characteristics', in E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, <plato.stanford.edu>.
- Haig, D. and Hurst, L. D. [1991]: 'A Quantitative Measure of Error Minimization in the Genetic Code', *Journal of Molecular Evolution*, **33**, pp. 412–7.

- Hazel, W., Smock, R. and Lively, C. M. [2004]: 'The Ecological Genetics of Conditional Strategies', *American Naturalist*, **163**, pp. 888–900.
- Jablonka, E. [2002]: 'Information: Its Interpretation, its Inheritance, and its Sharing', *Philosophy of Science*, **69**, pp. 578–605.
- Jablonka, E. and Lamb, M. J. [2005]: *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*, Cambridge, MA: MIT Press.
- Janz, N., Söderlind, L. and Nylin, S. [2009]: 'No Effect of Larval Experience on Adult Host Preferences in *Polygonia c-Album* (Lepidoptera: Nymphalidae): On the Persistence of Hopkins' Host Selection Principle', *Ecological Entomology*, **34**, pp. 50–7.
- Lachmann, M., Sella, G. and Jablonka, E. [2000]: 'On the Advantages of Information Sharing', *Proceedings of the Royal Society of London B*, **267**, pp. 1287–93.
- Lehrman, D. S. [1953]: 'A Critique of Konrad Lorenz's Theory of Instinctive Behavior', *The Quarterly Review of Biology*, **28**, pp. 337–63.
- Lehrman, D. S. [1970]: 'Semantic and Conceptual Issues in the Nature-Nurture Problem', in L. Aronson, E. Tobach, D. S. Lehrman and J. Rosenblatt (eds), *Development and Evolution of Behavior*, San Francisco: Freeman.
- Leimar, O. [2009]: 'Environmental and Genetic Cues in the Evolution of Phenotypic Polymorphism', *Evolutionary Ecology*, **23**, pp. 125–35.
- Leimar, O., Hammerstein, P. and Van Dooren, T. J. M. [2006]: 'A New Perspective on Developmental Plasticity and the Principles of Adaptive Morph Determination', *The American Naturalist*, **167**, pp. 367–76.
- Lewontin, R. C. [1978]: 'Adaptation', *Scientific American*, **239**, pp. 156–69.
- Levy, A. [2011]: 'Information in Biology: A Fictionalist Account', *Nous*, **4**, pp. 640–57.
- Lind, M. I. and Johansson, F. [2007]: 'The Degree of Adaptive Phenotypic Plasticity is Correlated with the Spatial Environmental Heterogeneity Experienced by Island Populations of *Rana temporaria*', *Journal of Evolutionary Biology*, **20**, pp. 1288–97.
- Lorenz, K. [1965]: *Evolution and Modification of Behavior*, Chicago, IL: University of Chicago Press.
- Mameli, M. and Bateson, P. [2011]: 'An Evaluation of the Concept of Innateness', *Philosophical Transactions of the Royal Society B*, **366**, pp. 436–43.
- Maynard Smith, J. [2000]: 'The Concept of Information in Biology', *Philosophy of Science*, **67**, pp. 177–94.
- Maynard Smith, J. and Szathmáry, E. [1995]: *The Major Transitions in Evolution*, Oxford: Freeman.
- Millikan, R. G. [1984]: *Language, Thought and Other Biological Categories*, Cambridge, MA: MIT Press.
- Nygren, G. H., Bergström, A. and Nylin, S. [2008]: 'Latitudinal Body Size Clines in the Butterfly *Polyommatus icarus* are Shaped by Gene-Environment Interactions', *Journal of Insect Science*, **8**, p. 47.
- Oyama, S. [1985]: *The Ontogeny of Information: Developmental Systems and Evolution*, Cambridge: Cambridge University Press.
- Papineau, D. [1987]: *Reality and Representation*, Oxford: Blackwell.

- Pen, I., Uller, T., Feldmeyer, B., Harts, A., While, G. M. and Wapstra, E. [2010]: 'Climate Driven Population Divergence in Sex-Determining Systems', *Nature*, **468**, pp. 436–39.
- Sarkar, S. [2004]: 'Genes Encode Information for Phenotypic Traits', in C. Hitchcock (ed.), *Contemporary Debates in Philosophy of Science*, London: Blackwell, pp. 259–72.
- Shea, N. [2007a]: 'Representation in the Genome, and in Other Inheritance Systems', *Biology and Philosophy*, **22**, pp. 313–31.
- Shea, N. [2007b]: 'Consumers Need Information: Supplementing Teleosemantics with an Input Condition', *Philosophy and Phenomenological Research*, **75**, pp. 404–35.
- Shea, N. [2011a]: 'What's Transmitted? Inherited Information', *Biology and Philosophy*, **26**, pp. 183–189.
- Shea, N. [2011b]: 'Developmental Systems Theory Formulated as a Claim About Inherited Information', *Philosophy of Science*, **78**, pp. 60–82.
- Shea, N. [forthcoming a]: 'Cue, Signal, Inherited Information', in B. Calcott, R. Joyce and K. Sterelny (eds), *Signaling, Commitment, and Emotion*. MIT Press.
- Shea, N. [forthcoming b]: 'Genetic Representation Explains the Cluster of Innateness-Related Properties', *Mind & Language*.
- Shea, N., Pen, I. and Uller, T. [2011]: 'Three Epigenetic Information Channels and their Different Roles in Evolution', *Journal of Evolutionary Biology*, **24**, pp. 1178–87.
- Stegmann, U. E. [2009]: 'DNA, Inference, and Information', *British Journal for the Philosophy of Science*, **60**, pp. 1–17.
- Sterelny, K., Smith, K. C. and Dickson, M. [1996]: 'The Extended Replicator', *Biology and Philosophy*, **11**, pp. 377–403.
- Sultan, S. E. [2000]: 'Phenotypic Plasticity for Plant Development, Function, and Life History', *Trends in Plant Science*, **5**, pp. 537–42.
- Sultan, S. E. and Spencer, H. G. [2002]: 'Metapopulation Structure Favors Plasticity Over Local Adaptation', *American Naturalist*, **160**, pp. 271–83.
- Suzuki, Y. and Nijhout, H. F. [2006]: 'Evolution of a Polyphenism by Genetic Assimilation', *Science*, **311**, pp. 650–2.
- Uller, T. [2008]: 'Developmental Plasticity and the Evolution of Parental Effects', *Trends in Ecology and Evolution*, **23**, pp. 432–8.
- Uller, T. and Pen, I. [2011]: 'A Theoretical Model of the Evolution of Maternal Effects under Parent-Offspring Conflict', *Evolution*, **65**, pp. 2075–84.
- Waddington, C. H. [1942]: 'The Canalization of Development and the Inheritance of Acquired Characters', *Nature*, **150**, p. 563.
- Warner, D. A., Uller, T. and Shine, R. [2009]: 'Fitness Effects of the Timing of Hatching may Drive the Evolution of Temperature-Dependent Sex Determination in Short-Lived Lizards', *Evolutionary Ecology*, **23**, pp. 281–94.
- West-Eberhard, M. J. [2003]: *Developmental Plasticity and Evolution*, Oxford: Oxford University Press.