



COMPUTER SCIENCE: What Do Robots Dream Of?

Christoph Adami, *et al.*
Science **314**, 1093 (2006);
DOI: 10.1126/science.1135929

***The following resources related to this article are available online at
www.sciencemag.org (this information is current as of November 18, 2006):***

Updated information and services, including high-resolution figures, can be found in the online version of this article at:

<http://www.sciencemag.org/cgi/content/full/314/5802/1093>

A list of selected additional articles on the Science Web sites **related to this article** can be found at:

<http://www.sciencemag.org/cgi/content/full/314/5802/1093#related-content>

This article appears in the following **subject collections**:

Computers, Mathematics

http://www.sciencemag.org/cgi/collection/comp_math

Information about obtaining **reprints** of this article or about obtaining **permission to reproduce this article** in whole or in part can be found at:

<http://www.sciencemag.org/help/about/permissions.dtl>

This current understanding casts doubt on the long-standing assumption that the 9% volumetric expansion is critical for freezing to fracture rock, and the closely associated notion that the frequency and intensity of freeze-thaw cycles are the main environmental determinants of frost weathering. These cycles may nevertheless be important—not because they cross 0°C, but rather because large temperature gradients arise as the rock is brought into the critical subzero temperature range favorable for water migration and segregation-ice growth. These ideas pave the way to a more fundamental understanding of the effects of rock type and climate on frost weathering (8), and more generally, of the basic processes underlying many soils and landforms that are characteristic of cold regions, not only on Earth but also on Mars and other cold planets.

The realization that frost damage in porous materials is, in general, fundamentally related to water being driven thermodynamically into small cracks, where it forms segregation ice, provides a fresh perspective on other forms of rock breakdown. They include those due to moisture variations and salt crystallization at

above-freezing temperatures. As Taber (4) realized long ago, ice growth in soils is closely analogous to mineral crystallization in rocks. Modeling of frost weathering (1, 6) thus provides strong guidance for studying rock expansion and weathering due to wetting or the growth of salt crystals.

Fundamental insights into liquid water and freezing in confined spaces have recently emerged from studies of the premelting phenomenon (that is, the occurrence in most materials of liquid films on surfaces and interfaces at temperatures far below their bulk melting temperature). In a recent review, Dash *et al.* (5) discuss the physics of ice premelting and explore the diverse geophysical manifestations of the basic phenomenon on land, in the oceans, and throughout the atmosphere and biosphere. They discuss briefly the growth of segregation ice in rocks and the resulting fracture that were examined quantitatively by Murton *et al.*

Insights into ice premelting also have clear implications for various practical issues. They may lead to a better understanding of how concrete and other fabricated porous media degrade as a result of ice and salt

growth (9) and how to design more durable materials; such understanding remains elusive despite hundreds of publications on freezing in porous media (10). Premelting and freezing in confined spaces also have considerable relevance for the cryogenic preservation of organs, the cold storage of delicate foods, and the protection of stone monuments, buildings, and art work exposed to freezing conditions.

References

1. J. B. Murton, R. Peterson, J.-C. Ozouf, *Science* **314**, 1127 (2006).
2. K. Hall, *Earth Surf. Processes Landforms* **11**, 131 (1986).
3. G. P. Davidson, J. F. Nye, *Cold Reg. Sci. Technol.* **11**, 141 (1985).
4. S. Taber, *J. Geol.* **37**, 428 (1929).
5. J. G. Dash, A. W. Rempel, J. S. Wettlaufer, *Rev. Mod. Phys.* **78**, 695 (2006).
6. J. S. Walder, B. Hallet, *Geol. Soc. Am. Bull.* **96**, 336 (1985).
7. S. Akagawa, M. Fukuda, *Permafrost Periglacial Processes* **2**, 301 (1991).
8. B. Hallet, J. S. Walder, C. W. Stubbs, *Permafrost Periglacial Processes* **2**, 283 (1991).
9. G. W. Scherer, *Cem. Conc. Res.* **29**, 1347 (1999).
10. M. B. Snyder, D. J. Janssen, SHRP-C-203, *National Research Council* (1992).

10.1126/science.1135200


COMPUTER SCIENCE

What Do Robots Dream Of?

Christoph Adami

Perhaps robots aren't so different from us after all. Like us, they need to constantly ascertain where they are in the world, and like us, they work better if they have an accurate sense of self. On page 1118 of this issue, Bongard *et al.* (1) show that robots equipped with an algorithm that infers their own physical structure from stored sensory data—dreams of their prior actions, so to speak—perform better in a simple forward locomotion task than robots whose decisions are not dream-inspired. Furthermore, robots that use these self-models to plan future actions can recover autonomously from injuries, by adapting their gait to compensate for the changed circumstances.

A robot's most formidable enemy is an uncertain and changing environment. Typically, robots depend on internal maps (either



provided or learned), and sensory data to orient themselves with respect to that map and to update their location. If the environment is changing or noisy, the robot has to navigate under uncertainty, and constantly update the probabilities that a particular action will achieve a particular result. The situation

becomes even worse if the robot's own shape and configuration can change, that is, if its internal model becomes inaccurate. In most cases, such an event constitutes the end of that particular robot's adventure.

Bongard *et al.* aim to improve a robot's robustness in an environment that may include damage to the robot. At the beginning of a self-modeling cycle, a four-legged robot without an internal model of itself performs actions (while on a flat surface), and records its own response via tilt sensors and angle sensors in its

Robots that create and update internal models of their own structure may be able to navigate the world in a more robust way and provide a test bed for models of self-awareness.

joints. The robot then computationally tests candidate self-models, by re-imagining the actions it just performed and comparing the behavior of the model with its memory of the results—that is, the robot tries to explain the observed relationship between sensory data and leg actuation by making assumptions about its own configuration.

Even though the number of tested models is comparatively small (by only allowing a limited arrangement of limbs and their length), it is easy to imagine that many models can end up explaining the recorded behavior equally well (or equally badly). In the next stage of the cycle, the robot uses these equivalent models to find an action that would serve as the best way to discriminate among them. In other words, we could fancifully imagine the robot thinking: "Well, these three models all seem to work equally well with what I remember, but it seems to me that if I stick what I think is one of my legs out just so, then I can discover if I have a fourth leg or not." To narrow the choice of models, the robot then proceeds to test the action that provides the most information about the model's

identity in the real world, and the cycle begins again. After 16 such cycles, the robot tests the accuracy of the final self-model by performing a set of actions that, according to this model, will result in the largest linear distance traveled, and then executing these actions. The total distance traveled can then be used as a measure of the accuracy of the robot's model of itself.

An important feature of the cycle is the active role the robot plays in determining its best self-model. Bongard *et al.* tested this feature in control experiments in which the actions taken by the robot were not informed by the self-models: For example, they forced the action synthesis algorithm to simply return a random—rather than a maximally discriminative—action. Such passive strategies fared markedly worse, as measured by the actual distance traveled after the 16 cycles. But the most dramatic difference occurred when the length of one of the robot's legs was shortened after it had gained a good sense of self. In this case, the 16-cycle algorithm was run again, this time starting with the previous best model. The active algorithm enabled the robot to adjust its gait and regain forward motion, whereas the random action controls (that is, those in which the actions were not informed by the self-model) did not.

The algorithm used by Bongard *et al.*

makes use of key insights from information theory, namely, that minimizing entropy leads to maximum predictive power (2). A similar conclusion can be drawn for algorithms that strive to locate a robot within an unknown landscape: In this case, taking an active role in discovering the environment rather than solely relying on sensor data also leads to improved performance (3). Which leads us to wonder whether the approach of Bongard *et al.* could also be used to plan actions in a changing environment, based on modeling not of the self, but of the world. Active algorithms that use stochastic modeling of probabilities of beliefs (3) about the environment exist, but they cannot synthesize new environment models, nor generate appropriate behavior in them.

How would dream-inspired algorithms work in terra incognita? A robot would spend the day exploring part of the landscape, and perhaps be stymied by an obstacle. At night, the robot would replay its actions and infer a model of the environment. Armed with this model, it could think of—that is, synthesize—actions that would allow it to overcome the obstacle, perhaps trying out those in particular that would best allow it to understand the nature of the obstacle. Informally, then, the robot would dream up strategies for success—just as the robot constructed by Bongard *et al.*

dreams about its own shape and form—and approach the morning with fresh ideas.

Although such an algorithm would require far more complex simulations than those giving rise to self-models in the work of Bongard *et al.*, robots relying on this kind of navigation could play an interesting role in our quest to understand the nature of consciousness (4). For example, we ought to be able to record the changes in the robot's artificial brain as it establishes its beliefs and models about the world and itself, and from those infer not only its cognitive algorithms, but also witness the emergence of a personality. Thus, perhaps the discipline of experimental robot psychology is not too far off in the future. And even though the robots studied by Bongard *et al.* seem to prefer to dream about themselves rather than electric sheep, they just may have unwittingly helped us understand what dreams are for.

References

1. J. Bongard, V. Zykov, H. Lipson, *Science* **314**, 1118 (2006).
2. T. M. Cover, J. A. Thomas, *Elements of Information Theory* (Wiley, New York, 1991).
3. S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics* (MIT Press, Cambridge, MA, 2005).
4. C. Koch, *The Quest for Consciousness* (Roberts, Greenwood Village, CO, 2004).

10.1126/science.1135929

CELL BIOLOGY

Sara Splits the Signal

Juergen A. Knoblich

Extracellular signals, such as growth factors and hormones, are received by receptors at the cell's surface and then transmitted to the nucleus via distinct cascades of intracellular signaling molecules. Because many signaling molecules are associated with intracellular membrane-bound compartments, these compartments, the signaling components, and their activation states need to be equally distributed between daughter cells during cell division. This is particularly important in developing tissues, where morphogens can elicit concentration-dependent responses at very long ranges, and even small variations in their concentration can create very different effects. On page 1135 of this issue, Bökel *et al.* show that cells in the devel-

oping wing of the fly *Drosophila melanogaster* contain a specialized subset of intracellular vesicles called Sara endosomes, whose main function seems to be equally distributing components of the transforming growth factor- β (TGF- β) signaling pathway during cell division (1). This mechanism ensures that the activation state of the signaling pathway remains precisely the same in both daughter cells.

The effects of the morphogen TGF- β on vertebrate and invertebrate tissue development rely on a relatively simple pathway (see the figure). TGF- β binds to two cell surface proteins called type I and II receptors and induces their dimerization. The type II receptor phosphorylates and activates the type I receptor, which in turn phosphorylates the transcription factor R-Smad. Phosphorylated R-Smad binds to a co-Smad to form an active transcription factor that

translocates into the nucleus and induces the expression of target genes.

During mitosis, signaling molecules are internalized into specialized vesicles that associate with the mitotic spindle. This ensures equal distribution into daughter cells.

translocates into the nucleus and induces the expression of target genes.

Sara (Smad anchor for receptor activation), a conserved, membrane-associated adaptor protein, simultaneously binds to the TGF- β -receptor complex and the R-Smad (Mad in *Drosophila*) (2). Sara contains a so-called FYVE domain that binds phosphatidylinositol 3-phosphate [PI(3)P], a membrane phospholipid that is primarily found on early endosomes. Although earlier work suggests that Sara is required for TGF- β signaling and recruits the receptor-R-Smad complex into endocytic vesicles, subsequent reports have led to conflicting views on the precise function of Sara. Sara can also bind to the phosphatase PPIc—a negative regulator of TGF- β signaling—and therefore could also inhibit TGF- β signaling (3). Furthermore, experiments addressing the role of endocytosis in TGF- β signaling have given conflicting

The author is at the Institute of Molecular Biotechnology of the Austrian Academy of Sciences (IMBA), 1030 Vienna, Austria. E-mail: juergen.knoblich@imba.oeaw.ac.at