

David Harel proposes —

A TURING-LIKE TEST FOR MODELLING NATURE*

In 1950, Alan Turing proposed an ‘imitation game’, for determining whether a computer is intelligent. A human interrogator, the candidate computer and another human are put in separate rooms, connected electronically. Alice, the interrogator, doesn’t know which is the human and which is the computer, and has a fixed amount of time to determine their correct identities by addressing questions to them. The computer has to make its best effort to deceive Alice, giving the impression of being human, and is said to pass the Turing test if after the allotted time Alice doesn’t know which is which. Succeeding by guessing is avoided by administering the test several times. Here, I argue that a variant of this idea, but with a Popperian twist (Popper, 1959), is applicable to the computerised modelling of natural systems, particularly in systems biology.

1. The grand challenge

Many characteristics of man-made systems, especially those termed *reactive* by computer scientists, are central to the dynamics of biological systems too: heavy concurrency (simultaneity), event-driven and time-dependent behavior, cause-effect phenomena and distributed control. These occur from the atomic level, via the molecular and the intra- and inter-cellular levels, to full organisms and even entire populations, suggesting that biological systems can be modelled (that is, reverse-engineered, simulated and analyzed) using methods and tools developed for the engineering of complex man-made systems. Recent results from small-scale modelling efforts have been extremely encouraging, Harel (2003, 2005); Kam et al. (2003); Kugler et al. (2008); Popper (1959); Priami et al. (2001); Setty et al. (2008).

*This paper is a slightly expanded version of Harel (2005) and is published here with permission.

Most modelling efforts are partial, intended to capture some limited phenomena or mechanism, or a small part of a larger system. There is often a particular goal for the modelling, with specific laboratory observations relevant to it and specific behaviors for simulation and checking. The motivation for such modelling efforts is making carefully considered predictions, which would typically be driven by certain questions that the modeller has in mind. These predictions stemming from the model and its analysis then lead to particular experiments that are to be carried out in the laboratory, and thus to the discovery of new scientific facts. A different approach, far more ambitious in its scope and required work, is to aim at modelling a complete biological system, such as an organ or a full multi-cellular animal. This kind of effort is not motivated by specific questions but, rather, by the desire to understand an entire system. (See Harel (2003), in which I propose to model the *Caenorhabditis elegans* nematode.)

Such a comprehensive ‘grand challenge’ is extremely nontrivial and by its very nature is intended to take modelling to the limit: let’s model (reverse-engineer) a worm or an elephant similarly to the way we engineer a chemical plant or an F-15 fighter jet. The challenge is to construct a realistic model, true to all known facts, which is smoothly extendable as the facts are discovered. It would feature a three-dimensional, animated graphical front end and would enable interactive multilevel probe-able simulations of the animal’s development and behavior. The underlying computational framework would be mathematically rigorous but would also be intuitive enough for biologists to enter newly discovered data themselves. The model would also support alternative theses reflecting disagreement among the scientists, to observe and to compare their effects at run time.

As I have argued previously (Harel, 2003), achieving this goal, even for a very modest organism like *C. elegans*, would require enormous amounts of interdisciplinary work, both in the computational analysis realms and in the accumulation, assimilation and formalization of the biological data itself. It raises numerous difficulties, for some of which no solutions are known at present. The good news is that this is typical of many grand challenges, both past and future; like putting a man on the moon, proving Fermat’s last theorem or solving cancer. One of the characteristics of a ‘good’ long-term challenge is that, if successful, the results would be spectacular, but even if it is not completely successful, many fundamental and useful results will have been achieved along the way. In our case, a comprehensive *in silico* model of an entire living organism would constitute an unprecedented tool, allowing researchers to see and understand life in ways not otherwise possible, triggering and helping predict behavior, filling gaps and correcting errors, suggesting hitherto unimagined experiments and much more.

It is not my intention here to try to convince the reader of the virtues of such an effort, but many benefits can easily be imagined.

2. Measuring success

Still, what does being successful mean? How do we know when we are done, labelling the model valid? And are we ever done? It is one thing to build a computerised model that looks good and captures some desired, but limited, features of a biological system in order to provide answers to some specific research questions, but quite another thing to claim to have constructed a valid model of a full organism, using all that is known about it. In limited modelling, one has in mind a manageable set of laboratory observations – analogous to requirement sin engineering man-made systems – so that one essentially knows what to test for. The challenge in comprehensive modelling is so great and the required multi-levelled amounts of detail and their inter-combinations so vast, that it is no longer clear how to tell when one is done or what it means for one’s model to be valid.

To address this question, we must clarify what we mean by modelling a creature based on all that is known. We must decide upon the model’s levels of detail, so that we don’t find ourselves having to deal with quarks or quantum effects. Moreover, we cannot hope to ever find out everything there is to know about a complex organism, even after limiting the levels of discourse. A model enabling

computerised simulations can at best be based on the knowledge available at a given point in time and will have to take liberties with what we don't know yet or never will. For example, the model can be made to fill in certain dynamics by carrying out the biological equivalent of the movie industry's 'morphing' technique. In any case, this begs the question of how to tell when all that we do know has indeed been captured.

Here is where validating the model can be likened to the Turing test, but with a Popperian twist (Popper, 1959): a comprehensive model of a full organism will be deemed valid/complete/adequate if it cannot be distinguished from the real thing by an appropriate team of interrogators. This idea raises many subtle questions and may attract opposition on many grounds, which this short essay will not attempt to address. The reader is referred to Turing's original paper which discusses several issues that are relevant here, too.

3. Modifications to Turing's test

If we were to apply the idea in Turing's paper to validate models of natural systems, particularly biological ones, what types of modifications to the original test would we have to implement? First, to prevent us from using our senses to tell human from computer, Turing recommended employing separate rooms and electronic means for communication. In our version of the test, tailored for modelling a multi-cellular organism, we are not simulating communicable intelligence but development and behavior. Consequently, our 'protection buffer' will have to be quite more complex – intelligent, in fact! It would have to limit the interrogation (which would consist of probing the computerised model residing in one room and the actual *in vivo* or *in vitro* laboratory for the organism at hand that resides in the other) to be purely behavioural and to incorporate means for disguising the fact that the model is not an actual living entity. These would have to include neutral communication methods and similar-looking front-ends, as in Turing's original test, but also means for emulating the limitations of actual experimentation. A query requiring three weeks in a laboratory on the real thing would have to elicit a similarly realistic delay from the simulating model. Moreover, queries that cannot be addressed for real at all must be let unanswered by the model, too, even though the main reason for building models in the first place is to generate predictive and work-provoking responses even to those.

Second, our test is perpetually dynamic, in the good spirit of Popper's philosophy of science (Popper, 1959). A computer passing Turing's original test can be labelled intelligent once and for all because, even if we take into account the variability of intelligence among average humans, we don't expect the nature and scope of intelligence to change much over the years. In contrast, a model of a worm or a fly or an elephant that passes our version of the test can only be certified valid or complete for the present time. New research will repeatedly refute that completeness, and the model will have to be continuously strengthened to keep up with the advancement of science. The exciting point here is that once a model passes the test, that model can be viewed as a correct theory of the organism being modelled. In a sense it actually *is* the organism, unless and until it can be distinguished from the real thing by appropriately clever probes and the corresponding laboratory work, which is exactly what Popper is trying to teach us: the important work of experimental scientists is to refute accepted theories by increasingly sophisticated experimentation. And this is how science advances and knowledge increases. (By the way, the protection buffer required for our test will also have to change as advances are made in laboratory technology, but, interestingly, it will have to be made weaker, since probing the model and probing the real thing will become closer to each other.)

Third, our interrogators can't simply be any humans of average intelligence. Both they, and the buffer people responsible for 'running' the real organism in the room that contains the laboratory and providing its responses to probes, would have to be experts on the subject matter of the model, appropriately knowledgeable about its coverage and levels of detail. In the *C. elegans* proposal, for example, these would have to be knowledgeable members of the worm research community.

Clearly, the modified Turing test proposed here is not without its problems and is not put forward for immediate consideration in practice. Still, it could serve as an ultimate kind of certification for

the success of what appears to be a worthy long-term research effort. Variations of the idea are also applicable to efforts aimed at modelling and simulating other kinds of natural systems.

References

- Ciobanu, G., Rozenberg, G., (Eds.), 2004. *Modeling in Molecular Biology*, Springer-Verlag, Berlin.
- Efroni, S., Harel, D., Cohen, I.R., 2003. Towards rigorous comprehension of biological complexity: modeling, execution and visualization of thymic T cell maturation. *Genome Res.* 13, 2485–2497.
- Fisher, N.P., Hubbard, E.J.A., Stern, M.J., Harel, D., 2005. Computational insights into *C. elegans* vulval development. *Proc. Natl. Acad. Sci.* 102, 6, 1951–1956.
- Harel, D., 2003. A grand challenge for computing: full reactive modeling of a multi-cellular animal. *Bull. Eur. Assoc. Theor. Comput. Sci.* 81, 226–235.
- Harel, D., 2005. A Turing-like test for biological modeling. *Nature Biotechnol.* 23, 495–496.
- Kam, N., Harel, D., Kugler, H., Marelly, R., Pnueli, A., Hubbard, E.J.A., et al., 2003. Formal Modeling of *C. elegans* Development: A Scenario-Based Approach, Proceedings of the 1st International Workshop on Computational Methods in Systems Biology, Lecture Notes in Computer Science, vol. 2602, Springer-Verlag, Berlin, pp. 4–20.
- Kugler, N.K., Marelly, R., Appleby, L., Fisher, J., Pnueli, A., Harel, D., et al., 2008. A scenario-based approach to modeling development: a prototype model of *C. elegans* vulval cell fate specification. *Dev. Biol.* 323, 1–5.
- Popper, K.R., 1959. *The Logic of Scientific Discovery*, Hutchinson, London.
- Priami, C., Regev, A., Silverman, W., Shapiro, E., 2001. Application of stochastic process algebras to bioinformatics of molecular processes. *Inf. Process. Lett.* 80, 25–31.
- Setty, Y., Cohen, I.R., Dor, Y., Harel, D., 2008. Four-dimensional realistic modeling of pancreatic organogenesis. *Proc. Natl. Acad. Sci.* 105, 51, 20374–20379.
- Turing, A.M., 1950. Computing machinery and intelligence. *Mind* 59, 433–460.