David Wolpert Research Statement

Although my degrees were in physics (Princeton, then the Kavli Institute for Theoretical Physics), my interests are far broader. Reflecting this I have worked at the Santa Fe Institute, the Neurosciences Institute at Rockefeller University, Los Alamos' Center for Nonlinear Studies, TXN (a data-mining startup), and IBM. Currently I am a senior computer scientist at NASA.

Reflecting this breadth of my interests, I have published work on the following topics, most of it being on the first four:

Multi-agent systems and distributed control
Operations research (in particular nonlinear and heuristic optimization)
Machine learning
Statistics (Bayesian, Sampling theory, and Maxent)
Nanotechnology
Molecular biology
Computation theory
Game theory
Several branches of mathematics.

In the last few years my research focus has been five-fold:

1) PROBABILITY COLLECTIVES: By casting both of them in terms of information theory, game theory and statistical physics are, mathematically, identical. Intuitively, players in a game and their reward functions are identical to particles in a system and their energy functions.

This identity allows one to "cut and paste" techniques between these fields, creating a hybrid richer than either by itself. So for example, the proper quantification of the bounded rationality of an agent turns out to be formally identical to a particle's temperature. As another example, techniques in statistical physics used to analyze systems with varying numbers of particles can be applied to games with varying numbers of agents.

Along with the members of my group and collaborators at Oxford, Stanford, Berkeley, and Los Alamos, I have been investigating the extremely rich theory arising from this hybridization. This theory has allowed us to derive algorithms for distributed optimization in the broadest sense, and in particular for adaptive and robust control of multi-agent systems. In our own projects and in joint work with collaborators at GE and BAE, we have been exploring these algorithms on many real-world distributed control problems. So far these algorithms have always far outperformed conventional techniques in simulations and (a few) hardware demonstrations. Moreover, these gains always grow – often drastically – as the system size grows.

These real-world problems have come from telecommunications, adaptive programming of nanocomputers, dynamic rescheduling, control of mites on trailing edges of airplane wings to minimize turbulence, distributed design, and control of constellations of rovers.

2) SELF-DISSIMILARITY: With collaborators from the Santa Fe Institute, I have been working on measures for characterizing complex systems based solely on sample data from such a system, with no pre-existing theoretical model of the system. This work has been based on the insight that a defining feature of essentially any complex system (e.g., a city, a human body) is that the spatio-temporal patterns exhibited at different scales differ drastically from one another. Conversely, non-complex systems tend to be highly self-similar (e.g., a gas, a crystal, or a gas). This leads to the use of self-*dis*similarity measures to analyze complex systems. Intuitively, the self-similarity of a system is akin to the first moment of a distribution – it is all that is not captured in such a “first moment” that goes into defining the system’s self-dissimilarity.

Such measures can be viewed as quantifications of how different the computation processes at different scales are, and how information flows between the scales. Preliminary experiments with such measures have applied them to 1-dimensional bit strings and 2-dimensional visual images. With other colleagues we are now investigating their application to procedures in a hospital’s ICU and to financial time series data.

3) FUNDAMENTAL PHYSICAL LIMITS ON COMPUTATION: I am engaged in research on an alternative to the Chomsky hierarchy that is, I would argue, a better reflection of computation (and more generally observation, control, and other kinds of inference) as it actually arises in the physical universe.

This alternative hierarchy has a very rich mathematical structure, one that imposes a number of limits on what kind of computation, observation, and control can transpire in the physical universe. These limits do not rely on chaotic processes, limits on the Chomsky power of one’s computer, etc. In fact they are independent of the precise laws of physics governing our universe. As a result, they provide (for example) an uncertainty principle for observation that holds independent of quantum mechanics.

4) EXTENDING THE DEFINITION OF A METRIC: I have also recently worked on how to extend the conventional definition of a metric - intuitively, a measure of distance between two points - to measure “distance” within sets of more than two points. Amongst other issues, doing this has required extending the concept of the triangle inequality to involve more than two points.

This work promises to provide an algorithm mapping the characteristics of a space to an associated metric for comparing probability distributions over that space. No longer would one use a one-size-fits-all measure to quantify distance between probability distributions. This would potentially have applications throughout machine learning, statistics, and information theory.

5) PREDICTIVE GAME THEORY: Traditional noncooperative game theory hypothesizes that the joint mixed strategy of a game satisfies an “equilibrium concept”, with all other joint strategies being impossible. As an alternative one can (and arguably must) view the prediction of the joint strategy of a game as an exercise in decision theory. This means that one must first arrive at a probability distribution over possible joint mixed strategies. It is then the loss function of the external scientist making the prediction that picks which single joint strategy to actually predict; for the same game, different loss functions result in different “equilibrium concepts”.

Information theory and Bayesian reasoning can be used to arrive at the distribution over joint strategies. Among other things, the resulting mathematics provides a first-principles quantification of bounded rationality, one which explicitly arises as a cost of computation. With collaborators at Berkeley and Urbana-Champaign, I am using this quantification to measure how rationality changes in experimental human subjects under different conditions. Ultimately, the hope is to integrate this work with AI-style user-modeling, to generate a full-blown formalism for predicting the behavior of individual humans in economic scenarios.