

Modeling in Neuroscience as a Decision Process

P.R. Schrater (schrater@umn.edu)

Departments of Psychology and Computer Science, University of Minnesota
Minneapolis, MN 55455, USA

K.P. Körding (koerding@gmail.com)

Departments of Bioengineering and Neuroscience, University of Pennsylvania
Philadelphia, PA 19104, USA

G. Blohm (gunnar.blohm@queensu.ca)

Centre for Neuroscience Studies, Queen's University
Kingston, ON, K7L 3N6, Canada

Abstract:

Formal models play a crucial role in advancing knowledge about brain function. However, there are many incoherencies and questions typically asked in the field: what are good model goals? Which modeling tools are appropriate for a given question? When can we compare models? How can we know if model choices are appropriate and justified? This highlights some of the complexity of the modeling enterprise and suggests that modeling should be regarded as a high-dimensional decision process. Unfortunately, there is currently a lack of structure, principles and guidelines regarding this decision process. Here, we propose an organization of the problem, model and outcome spaces by describing modeling as a decision process that links model space to model utility via a crucial action space, i.e. model selection. We then provide decision heuristics and guidelines for best practices to help the modeler in the decision process. We hope this will help both modelers and reader better appreciate different models as an exercise of fruitful complementarity rather than a battle over what's the "right" model.

Keywords: meta-science; epistemology of modeling; model dimensions; utility; model outcomes

Introduction

What is the nature of the modeling decision? It begins when we detect a gap in our knowledge. Some observation or phenomenon reveals a gap in our knowledge. We notice something that provokes a question like "does that always happen", or "how does that work?", or "what causes that to happen?" For example in the moving train illusion, one might ask why do I perceive self-motion when I see the other train moving outside the window? For modeling decisions, the state space is comprised of the scientist's knowledge gaps. Scientists formulate knowledge gaps into problems, and we can think of the state as a point in a problem space. A modeling action is the choice of a mathematical representation that instantiates a hypothesized way to fill the knowledge gap. The outcome of modeling is to verify or refute that the hypotheses "work" in the sense that they can plausibly

fill in the gap, and this outcome has value according to the modeler's evaluation criteria.

We believe that formalizing the modeling process as a decision problem will provide the scientific community with clarity about their implicit goals, choices and preferences. We have previously proposed a guide on how to practically achieve this (Blohm et al., 2018). Below we flesh out the theoretical underpinnings of each of the decision components and argue that formalizing the decision problem turns the implicit and arcane art of modeling into an explicit choice problem that the scientific community can reflect on, improve and formally instruct. Specifically, this formalization suggests that modeling consists of a limited set of choices, i.e. defining the problem space (phenomenon, question, hypotheses, goals) and determining the desired model outcome space, i.e. what will I evaluate my model on? These choices will then naturally constrain the possible model space and thus simplify the model decision process. All these steps are detailed below and summarized in Figure 1.




Figure 1: modeling is a decision process. A specific question is selected from the problem space. Deciding about the modeling approach then allows formalization of the question and hypotheses and analysis of the results. Results are interpreted in the light of the initial question with respect to specific goals, which span the outcome space. Outcomes are evaluated to produce a utility measure.

The problem space

The problem space is what will define and constrain the actual modeling project and represents the knowledge gap. It is composed of four important and hierarchically linked components: (1) the **phenomenon** (P) at hand that requires better understanding, (2) the specific **question** (Q) that is of interest with respect to the phenomenon, (3) the detailed **hypotheses** (H) with respect to the research question, and (4) the particular modeling **goals** (G). We will outline each of these four dimensions and their linkages in this section, and we will illustrate these dimensions by using the well-known phenomenon of illusory self-motion induced by a moving train (see Figure 2).

Problem space

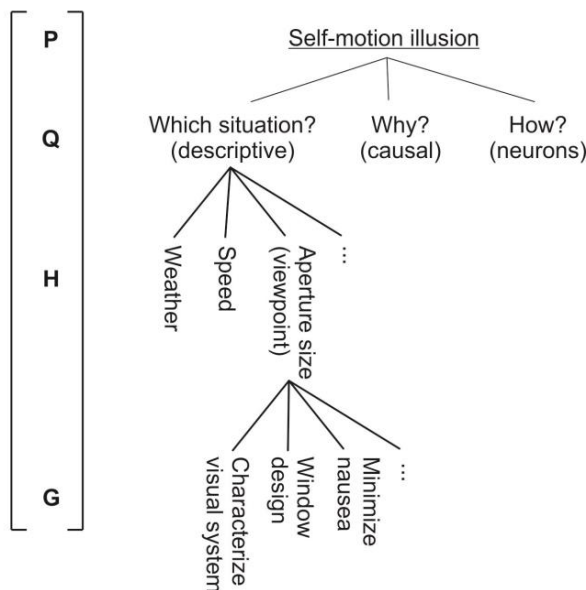


Figure 2: Example of problem space definition. A Phenomenon (P) can lead to different questions (Q). Each question can lead to different hypotheses (H) that are meaningful to achieve specific goals (G).

1. The **phenomenon** (P) at hand typically describes an experimental observation for which there is a lack of understanding. Ideally, a phenomenological description includes detailed observational data. Data richness and availability of different parametric studies can vary; the more data is available, the better the question, hypothesis and goal spaces can be constrained. As mentioned above, an example of such a phenomenon could be the train visual illusion of self-motion.

2. The lack of understanding of a phenomenon then triggers specific research **questions** (Q). Importantly, each phenomenon can potentially result in many different questions. For the train illusion, one might ask which parameters correlate with the strength of this

illusion, which would be a descriptive level question (see model space below). Alternatively, one might be interested in why this illusion arises, e.g. because of optimal integration of different conflicting sensory signals. Or a neurophysiologist might ask what neural properties underlie this illusion.

3. Once a question is formulated, researchers typically emit different **hypotheses** (H) about the expected relationships, mechanisms or properties underlying the phenomenon. These hypotheses crucially depend on the specific question asked. E.g., in the correlative question, one might hypothesize the set of independent variables that should play a role (e.g. speed, window aperture, weather conditions) or what functional relationship is expected (e.g. linear, quadratic, etc).

4. Typically, the modeler also has specific **goals** (G) of the model in mind that depends on the question and hypotheses. For example, in considering the aperture as a crucial parameter, one might be interested in optimal train window design (engineering goal), or in characterizing the visual system (basic science), or in reducing illusion-induced nausea (clinical).

This quartet of hierarchically linked components precisely defines the problem space. Of course there can be multiple hypotheses for a given question as well as multiple goals. Importantly, this description of the problem space defines the knowledge gap and we will use this below to inform the model selection process as well as the outcome evaluation. Indeed, as mentioned in Figure 1, we will propose a formal mapping between problem space, model space (this is what we will act on), outcome space and the utility of the outcome. We will do this in a probabilistic decision framework that allows us to optimize utility over all problem and model spaces. We argue that formalizing the problem space precisely will result in specific model selection (decision) heuristics.

Model dimensions and action space

Selecting a modeling approach is a decision process that is guided by the phenomenon we want to describe, the research question, our hypotheses and model goals. The problem space thus projects (maps) onto the model space. Here, we outline the different choice options, what model types exist and what questions they address.

Models are not equal. While there are many ways to classify models, we eschew classification by toolkit and instead focus on criteria that determine a model's breadth of utility. We argue that models vary along a few natural dimensions that determine well-defined limitations within which the model should be built and evaluated. We distinguish between two classification

criteria, an inferential criterion based on the generality of the model's application, and criteria based on the level of abstraction (granularity, scale). These criteria partially overlap but serve very different purposes.

Generality/granularity: The generality classification is based on the amount of structural knowledge (in terms of causal influences) present in the model, while the granularity (or abstraction) class is more concerned with how directly one can map a model onto detailed processes in the brain. Thus, each model has a different place in both classification schemes. Figure 3 depicts this combined classification scheme graphically with specific examples in each model category.

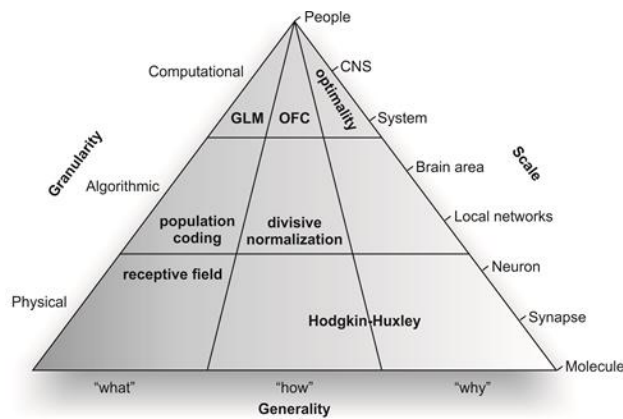


Figure 3: model classifications. Model generality (horizontal axis) is taken from Dayan & Abbott (2005). Granularity (left vertical axis) is adapted from David Marr's classification (1976). Model scale (right vertical axis) depicts the system scale and is adapted from Churchland & Sejnowski (1992). Different models types are shown as examples within each categorization. GLM: general linear models. OFC: optimal feedback control.

Abstraction: The vertical axes in Figure 3 depict the often-intertwined levels of granularity and scale of the physical elements modeled. These levels of abstraction form a hierarchy that is orthogonal to the generality (horizontal) dimension. Like most areas of science, models in neuroscience usually abstract out the details of levels of resolution finer than the phenomena of interest. The appropriate level of abstraction for a model is contentious, with individual researchers often having preferences for a particular scale of analysis. Authors have pointed out the value of abstraction and the importance of working at multiple levels of analysis but from different points of view.

Model outcome space and utility

Robert Rosen's book *Anticipatory Systems* (2012) describes the models as an attempt to bring a formal

system (the 'model') into congruence with a natural system. A goal is thus to capture the natural (causal) entailments of the natural system with the formal entailments of the model. Modeling is often regarded as an esoteric theoretical exercise with limited applicability to experimental research. While this unfortunately can be true, it is definitely not the goal. Rather, computational modeling of an experimental phenomenon has many tangible and important roles. It is well recognized that computational modeling is crucial for efficient and responsible research advancements. Here we list some of the unique insights gained through and advantages of computational modeling:

1. Representing causal linkages between brain and behaviour. Computational neuroscience is a rapidly growing field that can provide formal theories and frameworks to analyze and explain empirical findings using computational models. Importantly, it can bridge the gap between neural properties, computational objectives and behaviour. Thus, computational modeling is the only means by which one can attempt to obtain causal linkages between individual neurons or populations of neurons and behavior.

2. Guiding experimentation. One of the most important contributions of neurocomputation is to provide quantitative experimental predictions that can optimize experimental design, which is of particular importance when dealing with animal research or clinical patients. Smart computationally-driven experimentation can thus not only improve experimental outcomes but also be a guide in deciding which experiments are useful to do and which are not. This can be of particular importance when dealing with invasive methods requiring the sacrifice of animal life or when carrying out clinical investigations with limited tolerable patient access time.

3. Explicating complex phenomena. Computational neuroscience is also instrumental for our mechanistic understanding of brain (dys-) function. Formal models can provide quantitative mechanistic insight into complex system behavior that simple thought experiments often simply cannot. As a result, researchers can obtain a thorough understanding of a given experimental phenomenon.

4. Making assumptions explicit. The formalization of "word models" into mathematical language forces researcher to identify hidden assumptions and hypothesis underlying the model and missing knowledge.

5. Plan interventions by simulating changes in brain and behaviour. Building a theoretical model of the brain, of a neural system or even a neural

mechanism can be used to investigate dysfunction, e.g. after stroke. Such models of disease are paramount in devising novel treatments and rehabilitation strategies to improve the quality of life of patients.

6. Predict unobserved phenomena. Models can make targeted and testable predictions for experimentation. This is an important point. A model is a set of hypotheses that should be critically tested. The ultimate goal of this testing is its experimental validation/falsification, which in turn should lead to an improvement in the model.

7. Facilitate translation into applications. Finally, an often-neglected aspect of modeling is the fact that mechanisms by which the brain work can inspire new technologies. This might seem abstract for many researchers, but there are numerous examples where such additional valorization has occurred, e.g. attention models and Bayesian surprise used in surveillance or artificial networks used in stock market predictions.

Decision heuristics and best practices

We will use decision theory to clarify the modeling process. Decision theory provides a normative framework for making decisions by choosing an action that optimizes an objective given the current state of the decision-maker. Formally, taking an action a in a state s has a predicted outcome o with probability $P(o/s,a)$ and this outcome is evaluated according to a criterion function $U(o)$ called utility. Given a state s , we choose an action that maximizes the expected value of the criterion function. Formalizing modeling with decision theory means identifying the four concepts above - the state, action, outcome and utility functions.

With the above understanding of the model dimensions, we can now ask how to choose between all the possible combinations of approaches. We will argue that this is essentially an evidence-guided decision process. It is a decision process that is constraint (and often determined) by several pieces of information. (1) A specific research question often directly points at a subspace of models. (2) Specific hypotheses further constrain this space. (3) Explicit modeling goals further narrow down options. Sometimes, there might not be an existing modeling class that satisfies all the requirements; in this case the decision process might result in a compromise. Below we will unpack the individual pieces of information.

Let us consider the Outcome Space $OS = f(MS, PS)$ as a function of the model space (MS) and the problem space (PS) (see Figure 1). If $p \in PS$, $m \in MS$, and $o \in OS$, then we can express the utility of a model as a function of knowledge gained, i.e.

$$U(\Delta knowledge) = U(o) = \arg \max_{m \in M} U(f(m, p))$$

The decision here refers to selecting a specific modeling toolkit / methodology. We can then generalize our approach and ask what the expected utility of different problem spaces is, which can be written as

$$U(\Delta knowledge) = \arg \max_{m \in M} E[U(f(m, p))]_{p(P)}$$

We can thus formally treat the modeling exercise as a decision problem that can be optimized. Further details regarding the practical process of how-to-model can be found in our separate treatment of the topic (Blohm et al., 2018).

Discussion

Our goal was to provide a systematic approach to making meaningful decisions in modeling. To do so, we propose to frame the modeling endeavor as a decision process. This promotes identification of key model aspects and allows for the quantitative analysis of best modeling approaches for a given problem/goal set.

We believe that approaching the modeling exercise with the outlined rigor has key advantages in that it makes decisions, goals and desired utility of models explicit to the modeler and reader and thus increases transparency and clarity of models. We also hope that this will contribute to a more constructive and diverse set of models that are complementary, accepted by the research community and well justified in the field (Kording et al., 2018).

References

- Blohm G, Kording KP, Schrater P (2018). Beyond the Catwalk: A how-to-model guide for Neuroscience that dispels the mysteries. <https://osf.io/8qz65/>
- Churchland PS, Sejnowski TJ (1992). The Computational Brain. MIT Press
- Dayan P, Abbott LF (2005). Theoretical Neuroscience. The MIT Press
- Kording KP, Blohm G, Schrater P, Kay K (2018). Appreciating diversity of goals in computational neuroscience. <https://osf.io/3vy69/>
- Marr D, Poggio T (1976). From Understanding Computation to Understanding Neural Circuitry. AI Memo. MIT. <http://hdl.handle.net/1721.1/5782>
- Rosen R (2012). Anticipatory Systems. Springer-Verlag New York, Edn 2