*Design Patterns for Modeling and HCI [5/5]

(AKA #@! My Advisors Told Me)
Cognitive Modeling: A textbook on how to develop cognitive models

To appear in the Oxford Series on Cognitive Models and Architectures

Frank E. Ritter (frank.ritter@psu.edu)
College of Information Sciences and Technology (Applied computer science)
Penn State University
316g IST Building, Penn State, University Park, PA 16802
(814) 865-4453

The most perfect technique is that which is not noticed at all. -Pablo Casals, cellist, conductor, and composer (29 Dec 1876-1973)

*Alternative titles to show direction and intent

Design patterns for simulations (and HCI)
Design patterns for cognitive simulations
Applying unified theories of cognition: From architecture to furniture
Other potential titles: The Art of Modeling (Cognitive)
How to cognitive model
Popular Cognitive Modeling
Reusable design patterns for cognitive models
Reusable design patterns for models of cognitive mechanisms
Shit my advisors told me
Design patterns for modeling, simulations, cognitive science and HCI
Design patterns for modeling for cognitive science, HCI, and other sciences
Modeling design patterns
Modeling Cognition and Interaction: cognitive simulation
Popular Cognitive Modeling

*Preface [5/5]

There are methods of developing, testing, and presenting models that simulate cognition. When these methods are performed often enough or clearly enough, they can be seen as design patterns for how to create aspects of simulations and cognitive models of interaction. These models are used directly in science as theories of cognition and are used implicitly in design of interfaces.

These patterns are currently distributed across a wide literature. The space is not huge, but large compared to the time a single modeler has to do a class project, prepare a thesis, or work on funded research in a timely manner. These patterns keep getting rediscovered and reported by individual researchers, who have to do problem solving and research repeatedly to generate the most basic components of these types of simulation and modeling.

Also, many of these patterns are implicit, or even tacit. They are unwritten. The purpose of this book is to make these patterns more explicit, to write down some of the unwritten rules and strategies.

Having to start from scratch has impeded the use of simulations and cognitive models in psychology and other areas. This can be contrasted with the relatively fixed formula for running and presenting psychology experiments, which greatly helps with work in psychology. The format, the style, the order of psychology paper sections are all prescribed and proscribed for this type of work.

These patterns are also very related to work in human-computer interaction (HCI) and other simulation tasks. Projects to create models of cognition are often used in HCI, in that the HCI projects and researchers use models of users (sometimes explicitly and sometimes only implicitly) and use the same methods as modelers to build their model and to increase their understanding of the fit of the model of the user to actual user behavior. They use the results to improve the model of the user and to improve the design of systems, which is another goal of many cognitive modeling projects. So, this book notes and shows how these approaches are related in Chapter 1.

These methods are also developed and used by psychology where models of users inform the work, which should be every project in psychology. They should also be useful for anyone using or creating simulations.

Who the book is for

Thus, this book should be useful for a range of advanced undergraduate courses but will be used primarily by graduate students and researchers in universities, government, and industry. It will be primarily useful for those modeling human cognition, but it will also be helpful for those modeling other tasks, ranging from modeling societies to modeling sub-atomic particles.

There are design patterns for performing work across these domains that use models and theories to drive the work. Casting an idea into a formal representation is similar and has similar tasks and approaches. Developing the theory and realizing it as a set of equations or computer program has similar tasks and approaches. Preparing data to compare to the predictions and doing the comparison is similar and has similar approaches to make the work more productive. Presenting the theory, data, and comparison are similar and have similar methods. There are some aspects that are particular to generative models of cognition, but even these may find some

reuse in other areas of simulation and of science. So, this book should be useful to people using models, broadly defined.

It will not duplicate, but will be informed by books with some similarities, including *The engineering data compendium* (Boff, Kaufman, & Thomas, 1986; Boff & Lincoln, 1988), *Unified theories of cognition* (Newell, 1990), *How to lie with statistics* (Huff, 1973), *Design patterns* Gemma, 1995 #2026}, *Modeling human and organizational behavior: Application to military simulations* (Pew & Mavor, 1998) and the UK reply (Ritter et al., 2003), and works by Tufte (e.g., , 2006). These books provide theories and examples of how to provide methods and summaries of work in an area. They are extensively cited in the chapters.

What you will learn

This book will support several lessons about models and simulations.

The point of this type of work is insights, not a running model, not more data. This point has implications throughout the book. The book thus includes a description of what an insight is and numerous examples.

A description of the steps to create, test, and disseminate a model that will be helpful for cognitive modelers and other people working with simulations, interfaces, and theories, particularly formal theories.

What questions to ask. Knowing the right question to ask can be an insight and can greatly influence later work and may be the most important type of insight. For example, looking not just at how learning varies by trial but by time on task can open new results.

What data to gather. For example, not recording weight when giving a fixed dose of caffeine allows the participants' weights to add noise (caffein is not a dose dependent drug, but a dosage (dose/weight) dependent drug) when studying the effects of caffeine on cognition.

What factors to include in a study. A theory of how a task is performed makes strong predictions on what data should be recorded for later analyses. For example, working memory capacity, previous practice, and typing and general processing speed of the participants will influence task time. Time of flight and direction of particles will influence time through a particle detector. If all the features are not recorded, the data can often be useless for developing a better theory.

How to show that your model is worth taking seriously. You cannot prove your model (although there is a new stats technique that is helpful). You can, however, show that the model is worth taking seriously, and this is a complex mathematical and also pragmatic task, and has even social, cultural, and historical aspects with usability also being important.

Some of the relevant literature. The book includes numerous references. Read them if they are helpful; ignore them if they are not useful. There are often unexplored ideas in these papers and pointers to future work. Books that are reviews are particularly rich and often note future projects explicetly {e.g., \Pew, 1998 #1000;Pew, 2007 #2078}{Ritter, 2003 #1392}{Newell, 1972 #189}.

Approaches included

The book will include design patterns for performing common and powerful steps in modeling. These patterns are taken from influential papers or developed from sources where they are used. Where necessary, they have been created by the author and his lab.

Here are some example patterns:

- How to explain an architecture and model, structurally, functionally, and operationally
- How to present a simple graph of data
- How to automatically fit your model to data and why we increasingly need to do so.
- How to test a model of sequential actions: including how to reset your model after a mistake in a highly sequential task and keep score.

Running examples and previous work

The book will be informed by active projects to (a) model learning, (b) create an artificial ear (computational audio scene analysis), and the use of models to (c) improving situation awareness and (d) high-stakes test equipment, as well as numerous previous projects on modeling learning, behavioral moderators, problem solving, and interaction. The book is also informed by work running psychology experiments in learning and cognition, work on modeling learning and problem solving, and work modeling interaction and modeling behavioral moderators.

Finally, I introduce a theory of modeling and HCI in the first chapter to show how cognitive modeling, HCI (and human factors), and psychology and computer science are tied together. This theory provides a way to organize the book as well. It also shows how many of the methods are useful in both subfields to develop cognitive models per se in cognitive science, and to help HCI researchers with the aspects of their work that use user models (which is an enormous part of their work).

How the book will present material

This book will present common design patterns that will help—both novice and experts—to model more efficiently and clearly. The lessons will be drawn primarily but not exclusively from cognitive architectures communities. The book will support both general simulation work in behavioral modeling (which tends to mean interaction with the environment or movement through an environment) but will also support cognitive modeling (tending to model single user screen-based tasks). By cognitive modeling, I mean the modeling of cognition where the models of cognition process information to produce behavior. Where I say architecture, non-cognitive simulation folks might hear car and road, agent, or person in a social simulation. HCI and psychology folks can hear "my own mental model of what users are doing." These patterns apply quite broadly. Computer scientists can hear "my system".

This book is designed to be the first 12 to 15 hour lectures on this topic for graduate students or researchers moving into these areas. Thus, there will be topics further along the scope (like lecture 55, how to compare game scores to the power law of learning), and lectures too early (what is a t-test and how to perform it), that will not be included.

*Table of Contents [1/1]

***Forewords [1/1] 6**
***Preface [5/5] 12**
Who the book is for 14
What you will learn 15
Approaches included 16
Running examples and previous work 16
How the book will present material 16
***Table of Contents [1/1] 17**
***Ch. 1 Introduction [11/11] 21**
Abstract 21
Introduction 21
Why model? 22
What and what level to model 32
When to and why not build a model/ what level of model to build 34
The model will be wrong but may be useful and worth taking seriously 37
Thoughts on the general process and preview of the book 37
Comments on publishing whole models 47
Summary 49
Further readings on modeling in general 50
Exercises 53
Ch. 2 Choosing a modeling framework [3/9] 54
*Abstract 54
Introduction 54
*Types of models and what type to build 55
Choosing an particular framework 69
*If your architecture changes 74
Comments on publishing architectures 75
Summary 77
Further readings 78
Exercises 78
***Ch. 3 Choosing a task [12/12] 79**
Abstract 79
Introduction 79
A criterion-based approach to choosing a task 80
Usability of the task for the modeler 82
Recording human behavior on the task 83
Ways for models to interact with tasks 84
Common tasks 87
Creating your own task 90
Comments on publishing a task 90

Summary 91
Further readings 92
Exercises 92
Ch. 4 Choosing data to develop and test a model [8/13] 93
*Abstract 93
*Introduction 93
*Use a theory for gathering data 94
*Readings from psychology 94
*Readings about the domain 95
*Readings about the task 96
Types of data for building a model 98
Verbal protocols 8888 102
Running your own study 104
Publishing data for modeling 8888 106
Summary 108
*Further readings 108
*Exercises 109
Ch. 5 Building a model [7/13] 110
*Abstract 110
*Introduction: How to wrestle a bear 110
*Issues in building a model 111
*Patterns for creating models 119
Tools for creating models 123
Patterns in models 125
Patterns in inputs 127
Patterns in outputs 127
Patterns of learning 128
*Comments on publishing building a model 131
Summary 132
*Further readings 133
*Exercises 134
Ch. 6 Describing a model and its predictions [9/12] 135
*Abstract 135
*Introduction 135
*What audiences want to know about models 136
*Model traces and freezing your model 144
*Theory of how to create graphs and diagrams 144
*Example simple bar plot development 146
Example line graph comparing model and data 150
Complex graphs of behavior 150
*Comments on publishing a model description 151

Summary 153
*Further readings 154
*Exercises 154
Ch. 7 Evaluating your model [16/19] 155
*Abstract 155
*Introduction 155
*Overview: Worth taking seriously and can be improved 155
*Previous reviews on how to test models 163
*First, run your model until you understand it 168
*Scorecards as a model summary 171
*Fitting your model to data 174
*Product correspondence: Non-data methods to evaluate your model 178
Intermediate results: Using sequential data to evaluate your model 181
Temporal correspondence: Simple Methods to evaluate your model 183
*Error correspondence: Errors users do or could make 185
*Correspondence of context dependency 186
*Learning correspondence 187
*Other measures: Physiology, self-report, and intrinsic task scores 187
Advanced methods to evaluate your model 188
*Comments on publishing model evaluations 197
*Summary 199
*Further readings 200
*Exercises 200
***Ch. 8 Publishing your model [9/9] 202**
Abstract 202
Introduction 202
Places to publish modeling papers 203
The tactics of a modeling paper 208
The paper publication process 210
How to prepare a conference paper 219
Summary 221
Further readings 222
Exercises 223
Ch. 9 Example models and advanced issues [1/12] 224
*Abstract 224
Introduction 224
Example model 1: Higher-level programming language 224
Example model 2: Having rules compete 225
Example model 3: PDP model 225
Models as synthetic subjects, using models to test theories 226
Building large models: Software engineering and models 226

Ethics in modeling....227
Comments on publishing advanced issues....227
Summary....227
Further readings....227
Exercises....227
Ch. 10. Summary [4/8]....229
*Abstract....229
Introduction....229
Common misconceptions and arguments....**Error! Bookmark not defined.**
Future and current problems [drafting]....231
*Protecting your career as a modeler....236
*Summary....237
*Further readings....237
*Exercises....238
***References [1]....239**
***Author Bio [1]....257**
***Appendixes [3/3]....258**
*Appendix 1. General comments on writing....258
*Appendix 2. Commands to mark up a document....259
*Appendix 3. Publishers, journals, and conferences....260
References to absorb....265
Bones....267
***Glossary [1]....268**
***Index [1]....269**