

ALCINO J. SILVA
ANTHONY LANDRETH
JOHN BICKLE

ENGINEERING THE NEXT REVOLUTION IN NEUROSCIENCE

The New Science of Experiment Planning

OXFORD

ENGINEERING THE NEXT REVOLUTION
IN NEUROSCIENCE

This page intentionally left blank

ENGINEERING THE NEXT REVOLUTION IN NEUROSCIENCE

The New Science of Experiment Planning

Alcino J. Silva
Anthony Landreth
John Bickle

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trademark of Oxford University Press in the UK
and certain other countries.

Published in the United States of America by
Oxford University Press
198 Madison Avenue, New York, NY 10016

© Oxford University Press 2014

All rights reserved. No part of this publication may be reproduced, stored in a
retrieval system, or transmitted, in any form or by any means, without the prior
permission in writing of Oxford University Press, or as expressly permitted by law,
by license, or under terms agreed with the appropriate reproduction rights organization.
Inquiries concerning reproduction outside the scope of the above should be sent to the
Rights Department, Oxford University Press, at the address above.

You must not circulate this work in any other form
and you must impose this same condition on any acquirer.

Library of Congress Cataloging-in-Publication Data
Silva, Alcino Jose Roriz Teixeira da, 1961–

Engineering the next revolution in neuroscience : the new science of experiment planning /
Alcino J. Silva, Anthony Landreth, John Bickle.
pages cm

Includes bibliographical references and index.

ISBN 978-0-19-973175-6 (alk. paper) — ISBN 978-0-19-932904-5 —
ISBN 978-0-19-932905-2 1. Neurosciences—Research. 2. Neurology—
Technological innovations. I. Landreth, Anthony. II. Bickle, John. III. Title.

RC337.S55 2013
616.80072—dc23
2013005205

1 3 5 7 9 8 6 4 2

Printed in the United States of America
on acid-free paper

CONTENTS

<i>Acknowledgments</i>	<i>vii</i>
1. Navigating Neuroscience	1
2. Framing Maps of Research Findings	14
3. A Framework for Molecular and Cellular Cognition	26
4. The Power of Convergence	42
5. Ruling out the Probable to Argue for the Possible	75
6. Resonating Experiments	95
7. A Dream Within Reach: Building Automated Research Maps	117
8. Architecture for Experimentalists	138
9. The Future Science of Experiment Planning	161
<i>Notes</i>	<i>179</i>
<i>References</i>	<i>181</i>
<i>Index</i>	<i>187</i>

This page intentionally left blank

ACKNOWLEDGMENTS

From Alcino J. Silva

I would like to recognize the invaluable contributions of several people without whom this book would have never been possible. First, my wife Tawnie Silva for her endless patience, for her wisdom and advice, for her love and selflessness, for being there when there was little else; my parents, and my children Alex and Elenna, for their unwavering support, encouragement, and many personal sacrifices that allowed me to work on this book when I should have been with them; my friend Alan Smith for nurturing the seeds that grew into some of the ideas described here; and the young scientists that worked in my laboratory over the years—their talent, imagination, and razor sharp intellects were the fertile ground in which this body of work developed. Finally, this book would have never happened without my co-authors; John Bickle for taking my naïve enthusiasm seriously, and especially Anthony W. Landreth for the intellectual courage to dedicate 3 years of his life to this project. His ideas, vision and clarity transformed almost every page of this book.

ACKNOWLEDGMENTS

From Anthony Landreth

I would like to thank my wife, Kat Landreth, and my parents and grandparents, Christine and Mike Zimmerman and William and Shirley Clark, for their patience and support in the writing of this book. I would also like to thank the Carnegie Mellon Department of Philosophy for hosting me for a semester to learn the basics of graphical causal models. Special thanks go to Clark Glymour, Joe Ramsey, and Jim Bogen for conversations and lessons in the philosophy of science, and to George Graham for encouraging me to study philosophy of neuroscience in the first place. Thanks to the faculty at the University of Cincinnati's Philosophy and Sciences program. And, special thanks to John Bickle and Alcino Silva for the invitation to work on the project that yielded this book.

From John Bickle

I would like to thank my wife and lifetime partner, Marica Bernstein. Of all the things that life has thrown my way in a half-century on this planet, I'm most happy of all that the two of us are now living out our shared vision of joint human flourishing at our Farther Along Farm, amidst the red clay, rolling hills, and pine and hardwood forests of rural north central Mississippi. Thanks to three great step-daughters, Caroline Cooper, Kat Landreth, and Margaret Cooper, now all grown-up women, for enriching my life. Thanks to Mom and Dad, still going strong, in the closest life they'll ever get to "retirement." And many thanks, of course, to Alcino Silva and Tony Landreth, for putting up with my many worries about earlier accounts of these ideas. The initial ideas are Alcino's; the exceedingly hard work of putting them into initial written words was Tony's. My contributions were aimed at organization, presentation, idea development, and strength of arguments. This compelling

ACKNOWLEDGMENTS

project has now occupied us for more than a half-decade, ever since Alcino waxed philosophical one spring evening in 2005, sitting outside at the old Pavilion Restaurant atop Mt. Adams, overlooking downtown Cincinnati, the Ohio River, and the hills of northern Kentucky.

This page intentionally left blank

[1]

NAVIGATING NEUROSCIENCE

1. PLANNING EXPERIMENTS WITH A MAP AND COMPASS

Scientists often plan experiments using little more than their intuitions about what needs to be accomplished next. To the seasoned researcher, this procedure is akin to navigating a familiar but vast landscape without a map and compass. Scientists have been doing this for centuries now, since science's inception. This does not mean that experiments and research programs aren't often thoughtfully planned and executed. Scientists routinely think long and hard about which experiments to tackle next. But little prepares them to systematically scan the horizons of their field's potential. Good fortune and good mentors help some find successful routes among the many paths their research might take. Conversations with colleagues, journal clubs, and courses in experimental design offer some clues about how to steer research programs. But scientists won't find empirical research devoted to helping them decide how to structure a research program or what experiments to tackle next. In place of such research they will find only tradition, occasional lessons from their mentors, and their intuitions.

To extend our metaphor a bit: navigating over large, richly structured geographical areas without a map and compass is especially

problematic. And hopefully it won't take much argument to show that even the current body of published neuroscience research—the terrain that experimenters must now navigate to review and plan experiments—is overwhelming. Over the past three decades neuroscientists have published more than 1.6 million articles, spanning increasingly specialized yet increasingly complex and interconnected fields and subfields. Already this is a corpus far beyond human reading capacity. This terrain of knowledge is rich and highly structured, but with our current resources we have barely tapped it. How many unrecognized but important conclusions lay buried within the vastness of this record? How often do we duplicate published work, or pursue dead-end experimental paths that could have been avoided had we been fully aware of the contents and implications of the full published experimental record, even in our own fields? How much published research has actually advanced our understanding of nervous system functions? What percentage of this published research instead represents little more than small variations on previous findings, retracing of familiar ground in the immense landscape of the published record? If we had a tool to help us both navigate the vastness of the published record and grasp its implications, how might this tool change the course of current and foreseeable research? Even the most optimistic principal investigator, grant reviewer, or journal referee has to admit that these questions lack easy answers. Yet we can no longer afford to continue to conduct research at current scales—and expense—without addressing them. The corpus of neuroscience research against which we now ask these questions will only continue to expand.

Our purpose with this book is to address those who see that there is a real need to change the way we navigate the published record and plan future experiments. Our most earnest hope is that this book will convince readers of the necessity and urgency to develop

a science of experiment planning. The need for a systematic study of experiment planning is tightly connected with the urgent need for tools to navigate the immensity of the published record and inform experiment planning. Faced with increasingly complex research choices, neuroscientists now more than ever need new tools and new approaches to help guide their creativity and intuition.

2. A SCIENCE OF EXPERIMENT PLANNING

Statistics and the theory of experimental design enable us to answer formal questions about what can be inferred from our experimental data, given some assumptions about how the data were produced and how they are structured. But these disciplines operate at a level of abstraction far removed from many of the practical choices an experimenter must make when deciding how to produce that data. Experimentalists must either rely on existing paradigms or improvise when making these choices. There are no widely recognized principles to guide and inform experimenters' intuitions through the planning process, no systematic tools to help neuroscientists make these very important choices.

Before we go on and develop these themes, we will address a few questions that already may be troubling the reader. First, are we suggesting that we scientists don't plan our experiments thoughtfully? By no means. Do we doubt the power of scientific intuition? No, it would be silly to do so. Are we saying that we scientists possess no resources for systematically planning and evaluating experiments for relevance and importance? Again, by no means. Individual laboratories, grant review committees, journal referees, and research award committees think long and hard about whether

a proposed experiment has promise for the field or about whether results from a performed experiment deserve publication. These processes of evaluation are hardly unstructured, arbitrary, or capricious. Many of the evaluative concepts scientists use to make these thoughtful and often difficult decisions will be ones we'll appeal to in this book. But what we are claiming to be missing is a unifying framework, a scheme that consolidates these evaluative concepts, an organized effort to systematically study and evaluate strategies we currently use to plan experiments.

In conducting experiments, neuroscientists ask and answer “first-order” questions about nervous system phenomena and their causal interconnections (see Fig. 1.1). Examples include questions about the relation of a specific kinase to learning, or the expression of a specific gene to memory retention. The ideas that we are interested in exploring here concern instead a different class of questions—questions about research mapping and planning. Consider a first-order question about a given protein's relation to learning. About that we might ask: Can we determine the weight of the evidence supporting a role for this protein in learning? Based on this evidence, what are the different ways we might study that protein's relation to learning? Given the experiments that have been performed on this relation and the results so far gathered, what is the most relevant experiment to try next? These questions are “second-order”; they are questions about scientific practice. We will therefore follow common practice and call systematic, data-driven studies investigating second-order questions about science “S2 studies.”

To sharpen our statement of this book's guiding assumption: we propose that to answer S2 questions, we will need a framework of experimentation, a taxonomy, which will help us to organize both completed and planned experiments. Not only will we need

S2: a science that studies second-order questions.

Second-Order Question

What is the best way to study the relationship
between this kinase and learning?

·
·
·

First-Order Question

What is the relationship between this kinase and
learning?

Figure 1.1

a framework for the different kinds of information that different types of experiments can reveal, we will also need to be explicit about the rules used in combining the results from those experiments to infer neural mechanisms. We'll refer to this later process with the term "Integration"—that is, the process of combining results across distinct experiments. With a framework of experiments and the various methods of Integration stated clearly and explicitly, we propose to build maps of published research that will be used during experiment planning. The framework, the Integration principles, the maps of research findings we propose to derive from them, and their uses are at the heart of our discussion in this book.

We say that we will develop maps of research. This mapping metaphor suggests graphical or pictorial answers to S2 questions. Just as we might query a map about where we are geographically located relative to an unfamiliar landmark, so we might ask where specific lines of research fit among all of the other work that has been published on a particular scientific question. In later chapters of this book, we will consider how relying on a map of research findings makes it possible to determine where individual research efforts are positioned relative to other studies in a field. As the old

saying goes, and every scientist knows first-hand, pictures often communicate more effectively than words alone.

Constructing a map requires a legend, a system with which to sort features of the landscape into meaningful categories. Any such legend will only be useful if it accurately maps the terrain. Map builders thus need field knowledge to ensure that they are charting the terrain accurately. In our case, this means expertise in the field of neuroscience we wish to map. We must be careful not to assume that the specific strategies used to build maps for one field will apply to other fields. And yet, we should not lose sight of features shared among fields.

For example, all experimentalists value *reproducibility* of results and *convergent* evidence. Different fields of neuroscience will pursue reproducibility and convergence using different instruments and approaches. For example, a cognitive neuroscientist will want to know how consistently the same behavioral protocol yields the same neural activation pattern in neuroimaging experiments (i.e., reproducibility). That same researcher may also want to know whether the results from such neuroimaging experiments agree with the predictions of a systems-level computational model (i.e., convergence). A neuroscientist working in molecular and cellular cognition, for example, may want to know whether a particular mutation in gene X reliably affects memory in the same way across labs (i.e., reproducibility). That same researcher may also want to know whether that gene X is activated during memory (i.e., convergence). Within their individual fields, these two neuroscientists both look for evidence of reproducibility and convergence.

That we can look abstractly at two very different experimental traditions and see commonalities in research aims should not be too big of a surprise. But not all research is experimental. Although reproducibility and convergence are just as important in

computational neuroscience, the standards used to evaluate results in this field will include criteria that will look different from what we find in molecular and cellular cognition and cognitive neuroscience. Although computational models will often be held accountable to experimental data, and thereby inherit research strategies from experimental neuroscience, we should not confuse the project of mapping theories with the project of mapping experiments. Our focus in this book will be on mapping experimental research, because that is what is familiar to us. But, we should not forget the diversity of valuable research traditions outside of our framework.

3. GETTING CONCRETE

At an abstract level, all neuroscientists value reproducibility and convergence. But the strategies they use to pursue those value differ. We cannot be sure how well we are doing if we do not illustrate our approach with work from a specific area of research. How can we know whether our framework and principles will generalize if we cannot demonstrate that they apply to a single field?

So that we don't get lost in the clouds, we will illustrate our ideas as concretely as possible by discussing a number of case studies drawn from the published experimental record. To avoid errors of misinterpretation, we have kept the case studies discussed in this book within our scientific comfort zone. Our area of expertise is molecular and cellular studies of learning and memory. To put this point at the outset, with utmost clarity and explicitness: The system of experiment classification (i.e., the framework) and the methods of research Integration discussed here are derived from common implicit and explicit practices found in our subfield of neuroscience—molecular and cellular cognition (MCC for short).

The principles used in MCC reflect those used in other fields of biology where molecular and cellular approaches have had an impact, such as development, immunology, and cancer studies. We have not composed the framework or the principles we describe out of nothing. We simply made them explicit and propose to use them systematically for mapping and planning experiments in molecular and cellular cognition (Matynia et al., 2002). Thus before we delve further into the MCC framework and Integration principles, we will first introduce this relatively young neuroscience field.

The goal of MCC is to develop molecular and cellular explanations of cognitive phenomena. Neuroscientists working in MCC use a wide range of experimental tools, including molecular manipulations (e.g., gene targeting, viral vectors, pharmacology), cellular measures and manipulations (neuroanatomy, electrophysiology, optogenetics, cellular and circuit imaging), and a plethora of behavioral assays. The interdisciplinary experiments that characterize research in MCC reflect a large cross-section of approaches and techniques within current neuroscience. This makes our field an interesting place for the sort of S2 research we propose; potentially our S2 studies may be useful to experimenters beyond MCC. However, we expect that each major field in neuroscience—from molecular and cellular, to systems, to cognitive neuroscience—will need to develop their own strategies that reflect the implementation of fundamental research mapping principles such as reproducibility and convergence.

We will spend many pages of this book describing and illustrating our S2 ideas with specific experiments from MCC. Of course there are other experiments from MCC and from other fields of neuroscience that could be used as examples to illustrate our key concepts. We apologize in advance to our neuroscience colleagues for our admittedly biased illustrative choices. We simply used the

examples we know best. We welcome descriptions of experimental practice and results from other fields that illustrate our concepts, and perhaps add to, refine, or even replace some of them. We emphasize that our discussions of MCC research are not intended as ends in themselves. Their purpose is to illustrate a system for mapping and planning neuroscience research. We hope this book will show that this general system is a tool that could be used by any experimenter in any scientific field.

Clarification of our purpose is crucial at the outset. A tool that some regard as an aid may be regarded by others as a hindrance. This is especially true when that tool is supposed to help experiment planning—an area of acknowledged scientific creativity and genius. Our recommendations for a science of experiment planning are no more intended to dictate the experiments a scientist *must* perform than an accurate map dictates a traveler's choice of future destinations. Our hope, rather, is that knowing with more clarity where we are, where we could go, and the ways we could get there, we will make better experimental choices as we move science ahead. As with any other tool, proven usefulness is the ultimate criterion for judging its worth.

4. BUT HOW MUCH OF THIS IS GOING TO BE NEW?

Before we even take our first substantive step, we recognize that some readers will be concerned or even apprehensive about our project. We've already noted that we'll be employing some familiar terms to denote key concepts in our framework and subsequent results—for example, S2 and Integration. Our concerns with the “pragmatic rationality” of experiment planning are shared by many