Neuron
NeuroView

The Need for Research Maps to Navigate Published Work and Inform Experiment Planning

Anthony Landreth[1] and Alcino J. Silva[1,*]
[1]Departments of Neurobiology, Psychiatry, and Psychology, Integrative Center for Learning and Memory, Brain Research Institute, University of California, Los Angeles, Los Angeles, CA 90095-1761, USA
*Correspondence: silvaa@mednet.ucla.edu
http://dx.doi.org/10.1016/j.neuron.2013.07.024

The increasing volume, complexity, and interconnectedness of published studies in neuroscience make it difficult to determine what is known, what is uncertain, and how to contribute effectively to one's field. There is a pressing need to develop automated strategies to help researchers navigate the vastness of the published record. Simplified, interactive, and unbiased representations of previous findings (i.e., research maps) would be invaluable in preparing research surveys, in guiding experiment planning, and in evaluating research plans and contributions. Principles normally used in weighing research findings, including reproducibility and convergence, could be automated and incorporated into research maps. Here, we discuss a series of recent advances that are bringing us closer than ever to being able to derive systematic, comprehensive, but also interactive and user-friendly research maps. These maps could revolutionize the way we review the literature, plan experiments, and fund and publish science.

Introduction

The amount of published research in neuroscience has grown to be massive. The past three decades have accumulated more than 1.6 million articles alone. The rapid expansion of the published record has been accompanied by an unprecedented widening of the range of concepts, approaches, and techniques that individual neuroscientists are expected to be familiar with. The cutting edge of neuroscience is increasingly defined by studies demanding researchers in one area (e.g., molecular and cellular neuroscience) to have more than a passing familiarity with the tools, concepts, and literature of other areas (e.g., systems or behavioral neuroscience). As research relevant to a topic expands, it becomes increasingly more likely that researchers will be either overwhelmed or unaware of relevant results (or both). Consequently, there is a pressing need for new tools to help neuroscientists navigate the complexity and size of published information (Akil et al., 2011). There is an urgent need to develop research maps—simplified, interactive, and unbiased representations of research findings—not only to clarify what has been accomplished, but also to serve as guides in choosing what will be accomplished next.

The problem of mapping relevant research (i.e., determining the information directly relevant to a particular research topic) is closely related to the problem of experiment planning (i.e., conceiving and evaluating a potential series of future experiments). In choosing which experiment to perform next, we proceed with the hope that our knowledge and training will provide firm footing for a trek into unknown territory. But without research maps, we risk missing key information while planning new experiments. We also risk conducting redundant experiments. So, how can these research maps be built?

Recent technological developments bring us closer to developing research maps in three different ways. First, we can now build databases of unambiguous and concise representations of experiments and their results. Second, to assess the evidential weight in favor of hypotheses found among these representations, we can now automate familiar kinds of reasoning used in our respective fields to evaluate evidence. For example, reproducibility and convergence of research findings are two of the principles universally used in neuroscience to weight research findings. Reproducibility is the ability of an experimental finding to be replicated independently with identical or similar procedures. Convergence reflects the ability of very different experiments to point to a single conclusion. Quantitative measures of reproducibility and convergence could be used to weigh the evidence for embedded causal hypotheses in research maps (Figure 1). Third, we can now develop effective protocols for sharing these representations, so that we can combine knowledge across research communities.

Unambiguous and Concise Representations of Experiments

An important component of a "research map" is a database of research summaries and their results. This database could then be used to generate an interactive graphical summary (i.e., a literal map) of that research. The cartoon in Figure 1 illustrates the key steps used to create a research map, including the extraction of experiments and findings from the primary research literature (Figure 1A), the derivation of a database of those findings (Figure 1B), which is then used to derive an integrated graphical representation of those experiments (i.e., a research map; Figure 1C), and suggest causal hypotheses (Figure 1D). Just as a GPS map affords different levels of zoom, someone reading a research map would be able to survey a specific research area at different levels of resolution, from coarse summaries of findings (Figure 1C) to fine-grained accounts of experimental results. The primary

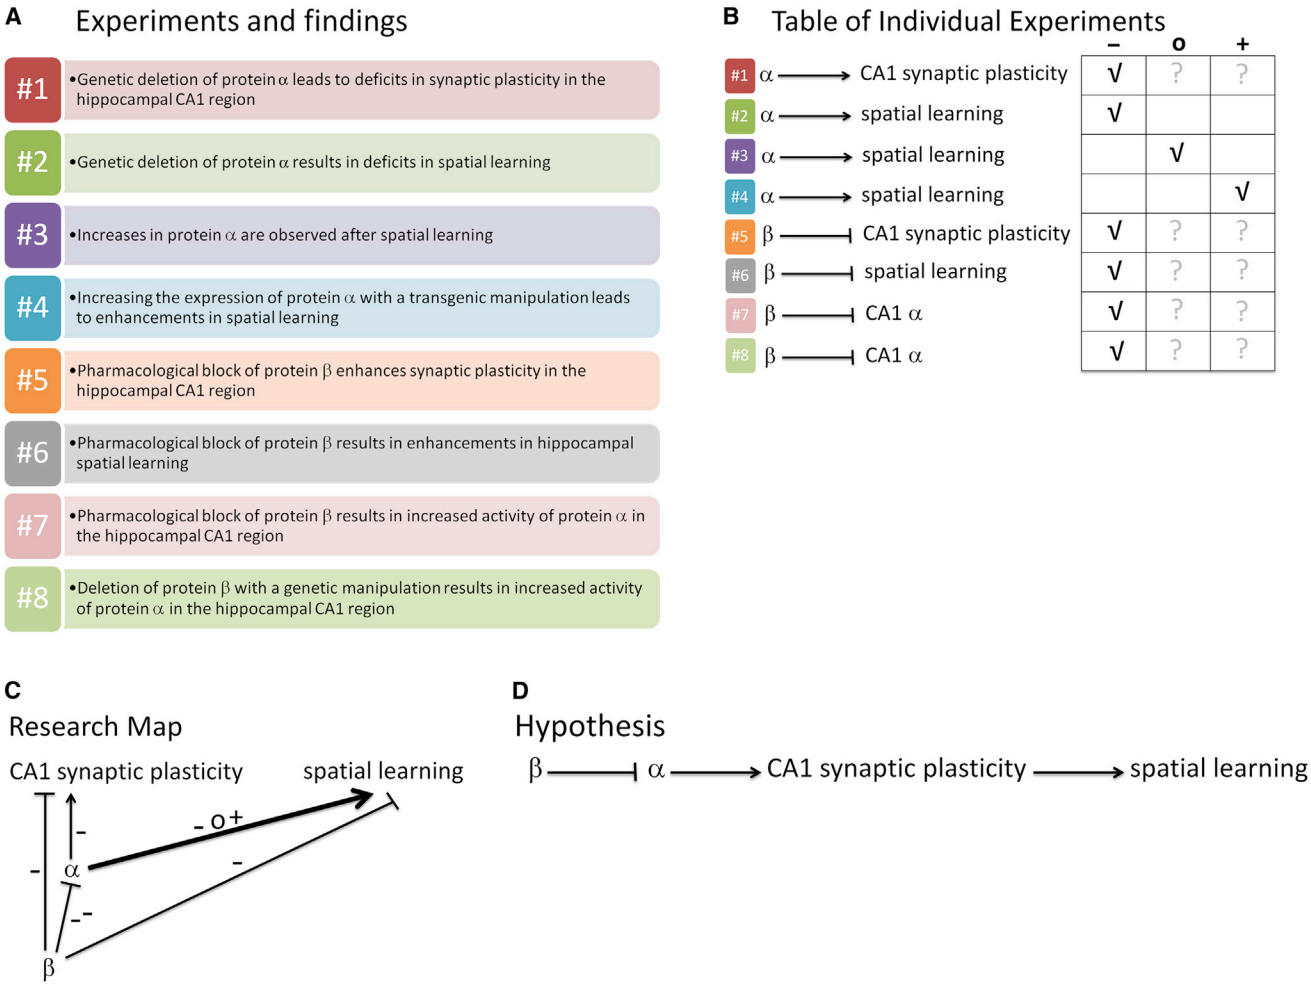

Figure 1. Steps Involved in Building a Neuroscience Research Map

(A) The first step in building a research map is to extract single experiments from a research publication. These individual experiments, along with other relevant information about methods and authors, could be captured into a nanopublication.

(B) The experiments could then be entered into a database with a format optimized for the extraction of graphic representations of those experiments (see below). The examples listed involve experiments with two variables (e.g., proteins β and α in experiment 8). The "−" symbol represents experiments in which the activity or levels of one variable were decreased and measurements were taken on another. The "+" symbol represents experiments in which the activity or levels of one variable were increased and measurements were taken on another. The "O" symbol represents experiments that involved no manipulation of either variable. Instead, the activity or levels of both variables were measured. The arrowheads in "+" and "−" experiments point away from the manipulated variables. In the "O" experiment, the arrowhead points away from the variable whose changes preceded changes in the other variable.

(C) A research map (integrated graphic representation) is derived from the database in (B). This map provides a convenient, although coarse, visual summary of the results listed in (A). The weight of the arrows represents the strength of the evidence supporting the proposed causal connection denoted by the arrow. For example, the connection with a heavy arrow is supported by the three types of convergent evidence outlined in (B), while the other connections with lighter arrows are supported by weaker evidence.

(D) Beyond providing objective summaries of experimental findings, research maps can also be used for hypothesis building. Depicted is a graphic representation of the hypothesis that β inhibits α and that α activation is needed for triggering synaptic plasticity in CA1, which in turn is required for spatial learning.

function of a research map is to display no more and no less information to a user than is necessary for the researcher's purposes.

Primary research articles often contain summaries of prior research and statements concerning the significance of findings presented. Additionally, review articles can help to place specific collections of findings in a broader and more integrated perspective. However valuable they may be, the individual perspectives in research papers and review articles are not always objective and balanced. Frequently, they do not reflect all of the relevant information available for the topic being reviewed. Thus, in addition to these personal perspectives, it would be useful to consult exhaustive, inclusive, and integrated databases (i.e., research maps) concerning the results and experimental strategies of an area or topic of interest.

To enhance the accessibility of research maps, each assertion would be stated in an unambiguous vocabulary. There are now numerous such vocabularies for automated reasoning, called ontologies (e.g., available through the National Center for Biomedical Ontologies, or NCBO). Unlike natural languages

(e.g., English), biomedical ontologies map one entity into one term. For instance, the word ''nucleus'' is ambiguous and could mean a cluster of cells, the nucleus of a single cell, and an atomic nucleus. The different senses of ''nucleus'' receive different terms in biomedical ontologies, so that when data are annotated with one of these terms, there is no ambiguity to confound a search over that data and no ambiguity to confound automated reasoning.

To date, the most extensive effort toward developing an ontology for neuroscience has been undertaken by the Neuroscience Information Framework (NIF). The NIF has collected a dynamic lexicon of over 19,000 neuroscience terms to describe neural structures and functions. The lexicon is built from the NIF standard ontologies (NIFSTD) (Larson and Martone, 2009). To make these vocabularies available to nonspecialists, the NIF group has built a web app, NeuroLex, from which a user can easily find the right terms to describe a phenomenon or protocol.

Ontologies like the NIFSTD provide materials for composing unambiguous representations of neuroscience research in a format sometimes called ''nanopublication'' (Groth et al., 2010). A nanopublication is the smallest unit of publishable information that can be uniquely identified and attributed to its author(s). Each of the eight experiments in Figure 1A could be reported in a single conventional research paper, or in eight nanopublications. Nanopublications usually include a subject-predicate-object structure, e.g., gene α (subject) is linked to (predicate) protein β (object). Nanopublications also provide metadata concerning, for example, the experimental methods used, as well as information about the authors (cf. http://nanopub.org). Together, these components of a nanopublication tell us no more than what we need to know when we search for specific results in the published literature.

Automated Reasoning

Nanopublications are a promising basis for building research maps. But to determine the evidential standing of the assertions found in nanopublications, it is key to know how and whether those assertions fit together. For example, are the findings underlying those assertions reproducible? Are there different sets of experiments converging on similar conclusions? When we informally ask these questions while conducting a literature review, we develop an intuitive sense of the robustness of a result or finding. With that sense, we decide whether we should trust a hypothesis enough to plan future related experiments. To be useful, causal connections represented in research maps would be weighted according to principles, including reproducibility and convergence, that neuroscientists use to weigh evidence for findings in their respective fields. For example, in Figure 1 there are three fundamentally different types of experiments supporting the idea that protein α is involved in spatial learning, while there is less experimental support for other potential causal connections listed in that figure. Neuroscientists have greater confidence in findings when they converge across different kinds of experiments. Similarly, results reproduced by multiple related experiments are deemed more reliable. For example, experiments 7 and 8 in Figure 1 both resulted in increases in the activity of protein α despite different methods to disrupt protein β (pharmacology and genetics). Reproducibility and convergence could be used to weight the evidence represented in research maps, and this would help identify strong versus weak results. To accomplish this, however, we would need to first organize experiments into categories.

For example, some neuroscience experiments are designed to decrease the probability of an event's occurrence, such as an inhibitory drug administered to prevent a receptor's action, or a lesion induced to impair a brain region's function (e.g., experiments 1, 2, 5, 6, 7, and 8 in Figure 1). Such experiments help us to determine the necessity of a specific phenomenon for the occurrence of another. Other experiments are designed to trigger an event, such as the expression of a gene or the activation of a brain region (e.g., experiment 4 in Figure 1). These experiments inform us as to the sufficiency of an event relative to the occurrence of another. In another common type of neuroscience experiment, no variable is intentionally manipulated and the goal is simply to describe how two phenomena covary, such as the activity of two molecules or two brain regions (e.g., experiment 3 in Figure 1). Not surprisingly, when the results of these three very different types of experiments agree, neuroscientists usually place more weight on the underlying hypotheses than when the support is incomplete (based on one type of experiment) or when there are contradictions in the results. One could imagine codifying this process in research maps, so that at a glance we could see the connections in research maps with weak and strong evidence. For example, the connection with a heavy arrow in Figure 1C is supported by the three different kinds of convergent evidence outlined above, while the other connections represented with lighter arrows have weaker evidential support. Unfortunately, it is often difficult to discern from literature searches, involving hundreds of papers and thousands of experiments, the weight of evidence (degree of convergence and reproducibility) behind any one finding. Research maps could be a solution to this increasingly serious problem.

In an attempt to represent large bodies of complex information, researchers draw diagrams with arrows (i.e., path diagrams) that stand for causal connections between phenomena, such as interactions between signaling molecules, and neuroanatomical connections (e.g., Figure 1D). These diagrams are useful for organizing existing research and planning future experiments. But these representations have important limitations. First, they are essentially static representations that do not update as the knowledge base of experimental results changes. Second, these diagrams do not show all of the equally well-supported alternative models that fit the existing data. Third, they do not show the relative weight of the evidence supporting each of the causal connections represented (commonly drawn as arrows). Finally, these diagrams are almost always composed by a small number of authors, and they are rarely systematic or complete. While the corpus of articles contributing to a diagram's composition is explicit in the review's bibliography, that corpus is necessarily subject to sampling biases, since a small number of authors will only be able read so many articles, recall so many facts, and reason over so many variables. Nor

is there an attending protocol that could enable others to read the same articles and thereby derive the same diagrams. Research maps could address all of these limitations while keeping many of the features (e.g., simplicity) that make these diagrams attractive to neuroscientists.

Ideas and strategies from graphical causal modeling (Pearl, 2000; Spirtes et al., 2000) will be useful for generating research maps. For example, very recently, an algorithm was developed that enables a collection of causal models with overlapping variables to be integrated into a unified causal network (cf. Tillman et al., 2009), a critical step in the generation of integrated large-scale causal networks. Imagine, for example, the complexities of attempting to integrate many related research maps such as the one in Figure 1C. How could this be accomplished in a systematic and automatic manner? The algorithms in graphical causal modeling could help us construct these integrated research maps, and these maps could be dynamically updated as new results emerge in the research record.

With a dynamic and interactive graphical interface, a scientist could use a research map to survey a field's experimental findings far faster than by reading abstracts or other textual descriptions. Areas with little research investment would be made apparent by both the sparseness and weakness of connections among their phenomena, enabling researchers to easily identify opportunities to conduct complementary experiments (for example, the experiments marked by "?" in the table in Figure 1B).

Currently, contradictions in the literature are difficult to resolve. These contradictions, however, would be accounted for in research maps by weakening the affected causal connections. Additionally, the global perspective afforded by these maps may help neuroscientists identify the source of contradictions or inconsistencies in the experimental record (e.g., by identifying systematic methodological differences between experiments with contradictory results). Research maps may also help address more objectively the quality of the evidence in the research literature. The uneven quality of research contributions is a real problem in science. Research maps will not solve this problem, but because they include databases of the information associated with research findings (e.g., methods, authors, tools, and models used), they may provide strategies to identify systematic problems in the research record.

Research publications normally highlight only a small subset of the research findings described. Most published experiments are not even alluded to in the abstract, and many are relegated to supplemental figures. Sadly, all scientists know that most experiments are not published at all and lay forgotten in research notebooks. This large body of forgotten research could be reviewed, reported as nanopublications, and integrated into research maps. Traditional research papers have to face the limitations of page counts, numbers of allowed figures, the attention span of potential readers, etc. None of these limitations would apply to the nanopublication content of research maps.

Shared Representations

Conceptually, it is not difficult to understand how research maps could be constructed (see cartoon in Figure 1). As a practical enterprise, the challenge might seem more daunting. Training in biomedical ontologies is not a core skill among experimentalists. Nanopublications are not part of the mainstream publication process. Natural language processing systems cannot yet automate the process of reading research papers for us, much less derive automated databases and graphic representations of findings from these publications. Time limitations and tradition also make the prospects for collective participation in the research mapping enterprise unlikely. How then are we to build research maps?

We can presently identify at least three strategies for building research maps. These strategies are not mutually exclusive. The first is a publically funded data entry effort. Specialists in various fields of research could be hired to write nanopublications for papers in their field. The database of nanopublications could then be deployed with a graphical interface. Forums, where the research community could critique the process, would be critical for the development and quality control of this effort.

The second strategy for building research maps piggybacks on activities that are part of the research community's typical workflow, such as note taking. From the time that they are students to the time that they are principal investigators, researchers take notes on the papers that they read. Cloud-based note taking applications (e.g., Evernote) could be used to weight, integrate, and eventually share these notes. If the workflow for note taking took the form of nanopublications, papers could be transcribed into nanopublications as an automatic byproduct of researchers doing what they already do. For example, a question and answer workflow could be developed for an online PDF reader. As a user reads research articles, questions about experiments are asked and, when answered, yield a database of structured notes for the user (and everyone with access to that database). This database would be useful to the user, as a simplified record of what was read, and useful for generating research maps as well.

The third strategy for building research maps builds nanopublications into the existing publication process. Different approaches could be taken toward implementing this strategy. For example, Microsoft has developed a plugin that assists authors in using ontologies to markup their text as they write. The markup could be used to render future papers machine readable. This would be an indirect approach. A more direct approach would incorporate fields for nanopublications into the templates for journal article submission. The NCBO makes an autocomplete widget for such purposes freely available. The widget will recommend terms from NCBO-hosted ontologies when a user has started typing in a data entry form field. The nanopublications resulting from filling out these forms could be published to a public database, just as abstracts are published to PubMed. As illustrated in Figure 1, this type of database would be the starting material for the construction of research maps.

First Steps

It is no mystery why efforts to derive simplified representations of research findings have not gotten a lot of attention. We have had neither an explicit

framework nor a data infrastructure sufficient to make the approaches proposed here a cost-effective endeavor. Recent developments from neuroinformatics and machine learning can now help us to overcome these hurdles. There is a growing sense of urgency in neuroscience to formally address the problems of research planning and coordination (Insel et al., 2003). The time has finally come to build tools to both map previous findings and aid experiment planning. We hope funding organizations, such as the National Institutes of Health and the National Science Foundation, as well as private foundations, take on this cause. Even a token investment could have an enormous impact on catalyzing the intellectual and structural resources needed for building research maps for integrating and planning experiments. With their help, we could have interactive mapping and planning tools for biology in the next 10 years. In our experience, even tiny handmade maps like the one illustrated in Figure 1 have been useful in our research, since they helped us to entertain experiments and approaches that our intuitions had overlooked. We may one day look on the time of experiment planning before research maps with the same incredulity we reserve for the days when experimental analysis was done without the benefit of statistics.

REFERENCES

Akil, H., Martone, M.E., and Van Essen, D.C. (2011). Science *331*, 708–712.

Groth, P., Gibson, A., and Velterop, J. (2010). Inf. Serv. Use *30*, 51–56.

Insel, T.R., Volkow, N.D., Li, T.K., Battey, J.F., Jr., and Landis, S.C. (2003). PLoS Biol. *1*, E17.

Larson, S.D., and Martone, M.E. (2009). Front. Neurosci. *3*, 60–67.

Pearl, J. (2000). Causality: Models, Reasoning, and Inference (Cambridge: Cambridge University Press).

Spirtes, P., Glymour, C., and Scheines, R. (2000). Causation, Prediction, and Search, Second Edition (Cambridge: MIT Press).

Tillman, R.E., Danks, D., and Glymour, C. (2009). Integrating locally learned causal structures with overlapping variables. D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds. Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8–11, 2008 (Cambridge: MIT Press), pp. 1665–1672.