Progress Toward Openness, Transparency, and Reproducibility in Cognitive Neuroscience

Rick O. Gilmore^a, Michele T. Diaz^{a,b}, Brad A. Wyble^a, & Tal Yarkoni^c

^a Department of Psychology, The Pennsylvania State University

^b Social, Life, & Engineering Sciences Imaging Center

^c University of Texas at Austin

Corresponding author: Rick O. Gilmore, Ph.D., Associate Professor, Department of Psychology,
The Pennsylvania State University, University Park, PA 16802; rogilmore@psu.edu

Short title: Openness in Cognitive Neuroscience

Keywords: open science; reproducibility; data sharing

Abstract

Accumulating evidence suggests that many findings in psychological science and cognitive neuroscience may prove difficult to reproduce; statistical power in brain imaging studies is low, and has not improved recently; software errors in common analysis tools are common, and can go undetected for many years; and, a few large scale studies notwithstanding, open sharing of data, code, and materials remains the rare exception. At the same time, there is a renewed focus on reproducibility, transparency, and openness as essential core values in cognitive neuroscience. The emergence and rapid growth of data archives, meta-analytic tools, software pipelines, and research groups devoted to improved methodology reflects this new sensibility. We review evidence that the field has begun to embrace new open research practices, and illustrate how these can begin to address problems of reproducibility, statistical power, and transparency in ways that will ultimately accelerate discovery.

Introduction

Most cognitive neuroscientists seek answers to questions about what patterns of neural activity underlie perception, thinking, memory, and action, among other topics. In answering these questions, we marshal evidence from studies of human and animal behavior, nervous system structure and activity, the effects of endogenous and exogenous substances, patterns of disorder and disease, and trajectories of change across the lifespan. Our common aim is to reveal reliable, reproducible, and useful facts about the relationship between mind and brain. These facts depend crucially on the tools we deploy to collect and evaluate data and on how we report what we do or do not find. In this paper, we review the degree to which our field meets the scientific ideals of reproducibility, transparency, and openness.

Rigorous self-reflection and self-criticism about methodology have been core values in cognitive neuroscience for some time ¹⁻³. Efforts to foster widespread data sharing ⁴⁻⁶ and other open research practices have long histories. What strikes us as new and important enough to merit reviewing them in 2016 are developments that likely cheer the pessimist and the optimist alike. On the one hand, accumulating evidence suggests that many findings in psychological science may be difficult to reproduce ⁷; statistical power in brain imaging studies is low ⁸⁻¹⁰ and has not improved ¹¹ over time; software errors in analysis tools are common and can go undetected for many years ¹²; and a few large scale studies and databases notwithstanding ^{4, 13-14}, the open sharing of data, code, and materials is rare. On the other hand, we see a renewed focus on reaffirming reproducibility, transparency, and openness as essential core values in psychological science and related fields ^{7, 15-19}. This reinvigorated focus has begun to force greater clarity about what these values mean in practice ²⁰. We find that the emergence and

rapid growth of data archives, meta-analytic tools, software pipelines, and research groups devoted to improved methodology are genuine reasons for optimism about the future of an open, transparent, and reproducible cognitive neuroscience.

In the sections that follow, we discuss definitions of open science practices and why they might be important for the field. We then review some of the history of these practices, discuss a range of recent developments, and offer some speculations about what the near future might hold.

History of open science practices in cognitive neuroscience

What are open science practices? What does it mean to reproduce or replicate a study? Most researchers agree that discovering robust and generalizable findings is central to the scientific enterprise^{1-2, 16, 19, 21}, but what evidence determines success or failure in meeting the ideal? In a previous paper in NYAS, Bennett & Miller¹ sought to assess the reliability of fMRI results, and provided data about the diverse measures used to assess reliability in the functional neuroimaging literature of that time. In summarizing the results of a large sample of studies reporting measures of test/retest reliability, Bennett and Miller¹ observed that no agreement exists about what constitutes acceptable reliability, nor was there consensus on what measure or measures should be used to evaluate it. Half a decade later, Goodman and colleagues²⁰ argue that uncertainty and disagreement about the meaning of these concepts²² persists and that misunderstanding impedes progress toward solutions. In response, Goodman and colleagues suggest three new terms which we adopt here: methods reproducibility, results reproducibility, and inferential reproducibility. *Methods reproducibility* means that a different investigator is able to obtain the same results when applying the same tools and analytical

procedures used in a study to the same (i.e., original) dataset. *Results reproducibility* means that a new study with new data, collected following the original procedures as closely as possible, yields the same outcomes. *Inferential reproducibility* occurs when independent researchers come to similar conclusions about what patterns of data mean, based on their own replication study or a reanalysis of a prior study²⁰. For example, Goodman et al.²⁰ suggest that competing views about the implications of a recent high-profile study of replicability in psychology⁷ stem, at least in part, from a disagreement at this level^{23, 24}.

Clearly, to achieve methods reproducibility, research practices that accurately and precisely capture essential details about methods, data, and workflows must be deployed; to achieve results reproducibility, those elements must be made openly and freely available to the scientific community; and achieving inferential reproducibility requires, among other developments, the capacity to accumulate, analyze, and interpret large quantities of data²⁵⁻²⁶ in consistent ways. Thus, openness and transparency relate directly to reproducibility of all three kinds. Reflecting this sensibility, a diverse array of behavioral scientists have begun arguing that achieving the scientific ideals of a free and open exchange of information requires the widespread adoption of open and transparent communication practices^{15-16, 27-30}. How well has the field of cognitive neuroscience measured up to these ideals?

Methods reproducibility

Much of cognitive neuroscience research is computationally intensive, so the extent to which the field's methods are reproducible depends on whether complex computational workflows can be reliably regenerated. Whether measuring task- or non-task-related nervous system activity using EEG, fMRI, or brain structure using MRI, CT, or PET, cognitive neuroscience studies regularly

generate spatially and temporally dense data streams. Seemingly minor choices made at each step of an analysis pipeline—including experimental design, data acquisition, preprocessing, analysis, and reporting—can ramify and have important implications for reproducibility.

The complexity of the typical neuroimaging pipeline is visible from the earliest stages of data acquisition. While there are only three major manufacturers of MRI scanners, the machines run different pulse sequences, and even scanners from the same manufacturers do not often run the same software. At the preprocessing stage—even prior to statistical analysis—researchers face a bewildering array of options when considering when and how (or even whether) to account for subject movement, signal spikes, differences in brain anatomy, physiological confounds, and any number of standard concerns. Statistical analysis is no less complicated, as researchers must decide what kind of analyses to conduct (mass univariate, multivariate pattern classification, etc.), what search space to use (whole-brain, specific regions-of-interest), what statistical contrasts and multiple comparisons correction procedures to apply, and so on. The sheer magnitude of variation in analytical approaches underscores why computational reproducibility is so critical in cognitive neuroscience and why it has historically seemed so daunting. Put simply, without the ability to understand precisely what steps a research group took, it is doubtful that anyone else could ever reproduce the procedures.

In EEG, the diversity of methods is arguably even larger than in fMRI, with numerous manufacturers and a corresponding variety of technologies using different kinds of electrodes, amplifier settings, cap configurations, and software packages ³¹. There have been some efforts at standardization of analysis methods through the release of software packages, such as the Matlab-based EEGLAB ³² and ERPLAB ³³. These packages have the advantage of allowing

researchers to explore the data using a graphical interface while simultaneously generating an executable history script that records most of the analysis decisions. The BigEEG Consortium (www.bigeeg.org), an offshoot of the EEGLAB initiative, seeks to develop and promote data and metadata standards for EEG-based research that may eventually facilitate large-scale analysis and meta-analysis. But, by and large, EEG data collection and analysis involves equipment and workflows that vary considerably from lab to lab.

Of course, the complexity of neuroimaging data analysis is not itself the enemy. It is not the raw number of methodological and analytical choices *per se* that creates barriers to reproducibility; rather, the challenge lies in encoding those degrees of freedom in a standardized (and ideally, machine-readable) way. Fortunately, over time, the brain imaging community has converged on recommendations about what parameters should be reported and how³⁴⁻³⁵. Moreover, at least in fMRI, imaging data analysis software shows a significant degree of standardization. From the earliest days of human brain imaging, leading research groups in the U.S. and U.K. wrote and freely distributed analysis software. This led to the widespread adoption of common tools with similar, although not identical, algorithms. Concerns about the inferential consequences of using one tool over another have been largely alleviated by findings from Gold³⁶ and Morgan³⁷, and questions about the reliability of workflows using one or more tools have addressed by Strother and colleagues³⁸⁻³⁹ and others (but see^{2, 12, 40}). All of the major tools in common use—SPM, FSL, AFNI, and BrainVoyager—enable researchers to write scriptable workflows, built either on internal engines (BrainVoyager), widely available commercial (SPM-MATLAB), or free/open source software languages (Linux/unix shell, python, C/C++ for FSL and AFNI). Naturally, there are some important caveats to this seemingly rosy picture. One concern is that while existing software supports relatively standardized and highly processing workflows *in*

principle, whether researchers actually take advantage of those features *in practice* is a separate matter. The number of SPM, AFNI, and BrainVoyager users who commonly rely exclusively on automated scripting in their analysis workflows, as opposed to using more user-friendly, but inherently irreproducible, graphical interfaces, is not known. We speculate it is small. Moreover, even in labs that do conduct fully automated analyses, the sharing or publication of the corresponding scripts or data processing pipelines remains rare ⁴⁰. While the differences between pipelines can be subtle ⁴⁰, the margin for error is also small (many published results only barely survive statistical correction), so a lack of full reporting can severely impair reproducibility.

A second caveat is that perfect reproducibility may be impossible to achieve even when a researcher is armed with all of the original data and scripts used to generate an analysis. Operating system differences, untracked differences in implicit software dependencies, and other factors can sometimes produce numerical discrepancies that, while initially small, may magnify as they cascade through a workflow to the point of introducing qualitative differences in results ⁴¹. We discuss potential solutions to this problem (e.g., containerization) later; for present purposes, we note that acknowledging the intrinsic limits of methodological reproducibility does not grant researchers license to ignore best practices in automation and code sharing.

Importantly, brain imaging data are only part of the reproducibility story in cognitive neuroscience. It is also critical to understand how to reproduce the *psychological* components of cognitive neuroscience studies—most notably, the experimental design and its intended relationship to the latent constructs of interest. Here, the prospects for full reproducibility have historically seemed less promising. Most experimental tasks involve the presentation of

sequences of visual or auditory events and the collection of participants' behavioral responses -- button presses, mouse movements or clicks, vocalizations, or eye movements -- using computer programs that instantiate tasks custom-tailored by a research team to address particular questions of interest. There have been researcher-initiated efforts to develop controlled vocabularies that describe the range of cognitive tasks deployed in the literature ⁴²⁻⁴³. NIH has spearheaded the creation of a standard toolbox of easily deployable tasks ⁴⁴ and also the development of data repositories designed to capture metadata about behavioral tasks and their variants ⁴⁵ in standardized and searchable forms. However, such efforts notwithstanding, most cognitive neuroscience researchers employ customized tasks built using a variety of software and scripting environments (e.g., E-Prime, the Matlab-based Psychophysics Toolbox, PsychoPy, and DMDX). Tasks use customized image and sound components, and researchers rarely share the code, image, or sound files used in experimental tasks ⁴⁰. These practices limit the reproducibility of behavioral measures used in cognitive neuroscience and psychology as a whole ⁷. Of course, the rigid standardization of tasks and materials has its own significant flaws, including the possibility of stifling innovation and slowing progress. We suggest that more widespread and open sharing of behavioral tasks, code, and materials provides a constructive middle ground.

Results reproducibility

Assuming that independent researchers are able to reproduce the methods of one another's studies, how closely do the findings generated converge? The answer is: It depends. In principle, even differences as basic as the make of MRI scanner could undermine the ability to compare results across studies ⁴⁶, so considerable effort has gone into standardizing techniques that allow multi-site imaging studies to be carried out in rigorous and reproducible ways.

Fortunately, the viability of the basic technology is no longer in any serious doubt; abundant evidence demonstrates that all major brain imaging techniques are at least *capable* of producing highly convergent results across different sites and experimental procedures. Perhaps the best-known effort to demonstrate the basic robustness of results—not only in neuroimaging, but in other biomedical fields, is the Biomedical Informatics Research Network (BIRN, <https://www.nitrc.org/projects/birn/>). BIRN is a multi-site, collaborative research consortium that strives to advance understanding of brain research and brain disease through the principles of data sharing and collaboration. There were several different BIRN initiatives, including the morphology BIRN, the mouse BIRN, and the function BIRN (fBIRN), among others. Although the fBIRN's disease focus was schizophrenia, considerable effort went into developing generalizable models for multi-site data collection, best practices for research, and methods to facilitate the use of standardized processes across sites. One of fBIRN's biggest contributions was software that enabled the systematic investigation of how fMRI activation signals vary across sites, field strengths, and scanner platforms. The project also developed methods to control for these differences ⁴⁷. fBIRN scientists developed an automated Quality Assurance (QA) procedure based on a standard MRI phantom, and the team released freely available software that could be easily incorporated into any service center's data transfer pipeline ⁴⁸. The fBIRN also provided leadership in modeling inter-site reliability ⁴⁹: The same 18 participants were scanned in four different scanning sites. These analyses revealed that inter-subject variability was 10 times greater than inter-site variability; activation in many brain regions showed fair to good reliability; and measures of reliability increased with more runs of data.

More generally, the ability of techniques such as fMRI to produce robust and replicable findings is demonstrated by the rapid canonization of many initially surprising neuroimaging findings. For

example, the tendency of a spatially conserved frontoparietal “task-positive” brain network to increase activity when participants engage in effortful cognitive activity has been replicated so often with fMRI over the past two decades⁵⁰⁻⁵² that the result is now often treated as a *de facto* manipulation check in new experiments. Large-scale meta-analyses of hundreds or even thousands of neuroimaging studies at a time further demonstrate a marked degree of convergence on stable neural correlates for most major psychological processes, from pain perception to episodic memory to language production^{26, 53-54}.

Of course, it is one thing to establish that neuroimaging methods can consistently reveal broad mappings between cognitive processes and distributed brain networks, and quite another to establish that the specific pattern of findings generated by any single study can be reproduced with a high degree of fidelity in another study. Unfortunately, as previous commentators have observed^{1, 55}, it is unclear whether neuroimaging findings meet this criterion. Arguably, the central problem is not that results reproducibility is particularly low, but that it has been difficult to quantify, leaving open the question of how much faith one should put in the results of any given published study. We focus on two critical barriers to the comprehensive assessment of results reproducibility in cognitive neuroscience, emphasizing fMRI (though the same concerns apply to other commonly used methods such as EEG, MEG, and TMS).

A first major challenge is that careful comparison of results across independent sites and studies typically requires that the full results be openly shared between sites, yet initiatives promoting neuroimaging data sharing have historically met with limited success. An early pioneer in data sharing was the fMRI Data Center (fMRIDC) at Dartmouth College, founded in 1999^{6, 56-57}. Around the same time, several journals, most notably the *Journal of Cognitive*

Neuroscience (JOCN), tried to implement mandatory open data sharing—including deposition of raw data files (e.g., BOLD time series, anatomical images)—as a requirement for publication.

These efforts to foster increased transparency, while laudable, sparked controversy and backlash from the community. Opponents raised concerns about practical issues— technology, data formats, time and money constraints, and privacy—and cultural ones—the possibly negative impact of open sharing on individual scientific careers and advancement, questions about data ownership, and whether data sharing should be mandatory or optional ^{57, 58}. The backlash eventually led JOCN to backstep on the data sharing requirement, and when funding to maintain the archive ran out, fMRIDC stopped accepting new data. fMRIDC's architects argue that despite the setbacks, fMRIDC should be viewed as a successful pioneer in open fMRI data sharing ⁵⁶, whose experiences shaped the next generation of repositories like the 1000 Functional Connectomes Project (FCP; fcon_1000.projects.nitrc.org/) and its International Neuroimaging Data Initiative (INDI) ⁵⁸, the OpenfMRI project (openfmri.org) and NeuroVault (neurovault.org), the Human Connectome Project (humanconnectomeproject.org), and the NIMH-based National Database for Autism Research (NDAR; ndar.nih.gov). More broadly, fMRIDC helped fuel interest in, recognition of, and support for the essential role information infrastructure (neuroinformatics) plays in making widespread data sharing and reuse possible.

A second barrier to the evaluation of results reproducibility is a lack of consensus about what quantitative measures should be used ¹. One area of contention is whether the magnitude or spatial extent of task-related activation (or both) should be assessed. Measures of magnitude and spatial extent depend on criteria for determining which voxels are active, of course. Bennett and Miller ¹ argued that the plurality of measures of reliability reported *within* individual studies made it challenging to ask about the reliability of findings *across* studies. They reported that the

reliability of group-level results, using individual participants tested at different times, varied depending on the temporal gap between the tests, the specific tasks employed (sensory/motor vs. cognitive), design factors (block vs. event-related), the magnitude of activations, and other inter-individual factors. Importantly, Miller ⁵⁹ and others ⁶⁰ found that, like the difference between within- and between site variability found by the fBIRN team ⁴⁹, variability within participants across testing sessions was lower than variability between participants. Differences in tasks, the degree of selectivity of active voxels to those tasks, and subject motion appeared to be the biggest contributors to inter-subject variability ⁶¹. Nevertheless, Bennett and Miller ¹ noted that all of the studies reporting test/retest reliabilities had small sample sizes, foreshadowing concerns about limited statistical power that others raised in the intervening years ^{8, 11}.

In sum, we believe that, perhaps surprisingly, given the size of the primary literature, the jury is still out on the degree to which researchers should expect individual neuroimaging findings to replicate when repeated under similar conditions. There is little doubt that methods like fMRI *can* produce highly replicable results, and that many canonical findings are indeed highly robust; however, as the low-hanging fruit are plucked and researchers increasingly turn to subtler phenomena, it becomes more important for researchers to share data and results openly. Only in doing so can we progress toward consensus on criteria for evaluating the reproducibility of results.

Inferential Reproducibility

The challenge of generating reproducible inferences—where independent researchers come to similar conclusions about what patterns of data mean—has been a central concern in the field for many years. Numerous published reviews have highlighted conceptual and statistical

problems that threaten common neuroimaging inferences^{25, 62-63}. One source of concern is that the statistical power of most fMRI studies is well below conventionally adequate levels⁸⁻¹¹; in a recent review based on over 1,100 samples, Poldrack et al.⁶² found that the median fMRI study in 2015 was underpowered to detect anything but relatively large effects (Cohen's d of ~ 0.75) even when using relatively high-powered procedures (i.e., a one-sample t-test). This observation is worrisome not only because low power implies a high false negative rate and inefficient resource expenditure, but because it frequently leads to incorrect interpretations—the notion that effects are stronger and better-localized than they actually are⁶⁴⁻⁶⁵. This increases the false-positive rate⁸ across the literature as a whole.

A second set of concerns arises at the analysis stage. As Bennett et al.'s well-known “dead salmon” illustration showed⁶⁶, insufficiently stringent multiple corrections procedures can easily inflate the false positive rate—an observation echoed by numerous studies that have highlighted limitations with common correction methods^{12, 67-69}. Moreover, such analyses all assume a best-case scenario under which researchers are not (inadvertently) capitalizing on the many “researcher degrees of freedom” available in a typical fMRI pipeline⁴⁰. If one could formally account for “p-hacking” (i.e., data-dependent selection of analysis procedures), it is likely that the false positive rate would rise, perhaps substantially^{11, 70}.

Lastly, even if one sets aside the statistical issues involved in the *generation* of cognitive neuroscience findings and assumes for the sake of argument that most published findings are fundamentally sound, it does not follow that researchers will agree about how to *interpret* such findings. Indeed, trenchant concerns have been raised about some of the most common assumptions researchers make when interpreting neuroimaging results, ranging from basic questions about what the BOLD signal reflects to what kind of information is actually extracted from multivariate pattern analysis⁷¹⁻⁷⁴. Poldrack⁷⁵ has flagged the problem of *reverse inference*

as a particularly serious challenge, noting that the widespread approach of inferring mental function based on the pattern of observed brain activity results runs a high risk of failure—unless it is supported by an appropriate Bayesian analysis that directly estimates the probability of a given task or state occurring conditional on an observed pattern of activity that is based on a reasonable prior distribution ²⁶.

Of course, science is a difficult enterprise, and it is easy to find serious methodological or statistical problems with virtually any piece of scientific research. The key question is what steps are researchers taking to address inferential concerns and to ensure that research findings continue to improve in reliability over time. To this end, to consider more recent initiatives aimed at improving the reproducibility of cognitive neuroscience research.

Recent Initiatives

The focus on problems of reproducibility in scientific research as a whole has accelerated in the last several years ^{16, 19-20}, and its scope extends well beyond psychological and neural science. As a result, cognitive neuroscience is both a beneficiary of new tools that promise to improve reproducibility and a contributor to them. We show that, fortunately, our field has already begun to embrace new, open and transparent research practices that promise to mitigate or even eliminate many of the serious problems of methods, results, and inferential reproducibility.

Methods reproducibility

Concern about reproducible workflows and practices across the computational sciences has sharpened in similar ways ⁷⁶⁻⁷⁷. While the specific practices that make computations reproducible vary from one field to another, Sandve and colleagues ⁷⁸ summarize a set of steps that have broad applicability to cognitive neuroscientists. These include avoiding manual data

manipulation steps (using scripts not GUIs); keeping careful track of the provenance (history) of all data, including derived results; tracking versions of all software and data; and providing public access to all code, outputs, and data.

Several data analysis tools have been developed based on free open source languages like R (RStudio; rstudio.com) and Python (Jupyter; jupyter.org). These tools support the creation of interactive electronic notebooks that combine data manipulation and analysis code along with graphic visualizations and text-based commentary. The tools can be used with version control environments like git or mercurial, allowing the history of a project's data analysis to be captured. Version control software can be used to store and share software, analyses, manuscripts, and documents written in virtually any language (both human and computer). Coupled with web-based repositories like GitHub (github.com), BitBucket (bitbucket.org), or the Open Science Framework (OSF; osf.io), version control systems enable researchers to share the histories and current status of all project materials and data. Some researchers concerned about computational reproducibility have gone even further, by creating full software environments that can run a particular analysis and packaging them in a specialized “containerized” environment (e.g., Docker; www.docker.com) that can be distributed for others' use across a wide range of computer platforms. The use of electronic notebooks, version control software, and web-based open data repositories has begun to enable cognitive neuroscience researchers to produce open and transparent workflows that can be readily reproduced. The authors use many of these techniques in their own research workflows.

Other efforts focus on methods reproducibility *across* study teams. On such initiative is the development of the Brain Imaging Data Structure (BIDS; bids.neuroimaging.io), a new open

data format designed to facilitate the storage and sharing of data from brain imaging studies ⁷⁹. BIDS attempts to achieve an easily implementable file directory and data structure that captures critical data and metadata about brain imaging studies and some data about the behavioral tasks performed by participants. BIDS arose out of the work involved in creating the OpenfMRI (openfmri.org) data repository ³⁴, designed to allow researchers to openly share raw BOLD imaging data sets with sufficient information to permit re- or meta-analysis. The BigEEG project mentioned earlier represents a similar data format standardization initiative targeted at the EEG community.

On the data sharing side, modern platforms have picked up where pioneers like the fMRIDC left off, making it ever easier for researchers to distribute large neuroimaging datasets in a readily usable form. A major initiative focused on methods and results reproducibility is the Stanford Center for Reproducible Neuroscience (CRN; reproducibility.stanford.edu) formed in 2015 by Russell Poldrack and colleagues. CRN is developing data repositories for both raw neuroimaging datasets (an upcoming successor to the OpenfMRI platform) and whole brain statistical maps (NeuroVault.org) ⁸⁰. A long-term goal of the CRN is not only to facilitate sharing, but also to provide containerized, modular, and fully reproducible cloud-based tools that can be easily executed via a graphical web interface. This will bring reproducible state-of-the-art neuroimaging data analysis within reach of researchers who lack the resources to deploy their own pipelines locally ⁸¹.

One of us ⁸² has argued that many problems in reproducing the methods of behavioral studies could be ameliorated if *video* of all experimental procedures was more widely recorded and shared with researchers. Text-based methods sections with restrictive page or word limits

simply can't convey sufficiently detailed information about a study's methods so that it can be reproduced by another researcher. Sharing video can pose privacy risks, but the Databrary (databrary.org) digital library, a repository specialized for storing and sharing video, has developed a policy framework to share identifiable data with participant permission. Like OSF, Databrary has begun to serve as a web-based home for researchers to store and share data, metadata, and materials about the non-imaging-related portions of a study, including videos of experimental procedures, images, audio recordings, or displays. Databrary largely focuses on developmental and learning science research now, but may expand in the future.

In sum, the field is making rapid strides to improve the reproducibility of methods with the emergence of new tools, practices, centers, web-based data management systems, and data repositories.

Results reproducibility

Despite the acknowledged lack of consensus about how to measure and thereby evaluate results reproducibility¹ and the noted significant problems with statistical power, we find a number of encouraging developments concerning the reproducibility of results. Researchers continue to take seriously the effort to systematically measure the factors that influence test/retest reliability of responses across time and tasks, and these sorts of studies are increasingly common⁸³⁻⁸⁸. Other research programs focus on addressing questions about the long-term within-subject stability of responses⁸⁹⁻⁹⁰, and how the accurate assessment of within-participant differences might address questions about individual differences⁹³. There is increasing support for conducting and publishing the results of confirmatory studies⁹⁴⁻⁹⁵, thereby

rectifying some existing biases that often favor the publication of new, novel results over confirmatory ones.

Several large-scale cross-site imaging studies whose results were designed to be widely shared with the research community have been undertaken (e.g., the Human Connectome Project and the U.K. Biobank Project). Findings from these studies are beginning to appear ⁹⁵ with results that both confirm and extend current understanding. Perhaps equally important is the extent to which planning for the large-scale sharing of these sorts of data has led to publication of extensive details about processing pipelines ⁹⁶ and careful planning about how to make shared components useful to other researchers.

Policy makers and publishers have taken a renewed interest in how the context in which scientific research is conducted and results shared can influence reproducibility. The Consortium for Reliability and Reproducibility (CoRR) has developed best practice guidelines for the use the resting-state fMRI data available through the INDI archive ⁹⁶, and the Organization for Human Brain Mapping has created a Committee on Best Practice in Data Analysis and Sharing (COBIDAS) ⁹⁷. Following on the success of the ArXiv preprint service, increasing numbers of cognitive neuroscientists have begun to deposit article preprints in the BioRxiv preprint service (biorxiv.org/neuroscience). An effort specific to psychological science (PsyArxiv; osf.io/view/psyarxiv/) has begun with support from the Center for Open Science (COS; cos.io) and the newly formed Society for the Improvement of Psychological Science (SIPS; improvingpsych.org). High profile generalist and topic-specific journals are adopting data sharing requirements reminiscent of those JOCN attempted to implement 15 years ago; there are new journals, such as Nature Publishing's *Scientific Data*, focused on creating citable,

scholarly homes for well-curated datasets; and some journals (e.g., *Cortex*) have adopted a new publication format, the pre-registered report, that conducts a review of the methods and analysis plan prior to data collection in exchange for a commitment to publish the results regardless of the findings.

There have also been recent developments to improve appropriate usage and reporting of statistical tests by means of automated tools such as *statcheck* (statcheck.io)⁹⁸⁻⁹⁹, which looks for elementary errors in the reporting of individual statistical tests. A related tool, *P-curve* (www.p-curve.com)¹⁰⁰ uses the complete set of statistical results from a body of work to estimate the evidentiary strength in favor of a hypothesis. While largely focused on the psychological science literature, these initiatives bear close watching by cognitive neuroscientists as they illustrate how the standardization of reporting practices can lead to insights about the quality of research practices and strength or weakness of evidence across a broad published literature¹⁰⁰. In the case of *statcheck*, the system depends on the fact that most experimental psychology papers report statistical analyses in ways that allow pertinent parameters to be automatically extracted from the published texts. Clearly, the diversity of efforts focused on bolstering the reproducibility of cognitive neuroscience results have considerable forward momentum.

Inferential reproducibility

Since the 2010 Bennett and Miller review on replicability, new tools and practices that promise to bolster the reproducibility of inferences have been created and are being adopted at an accelerating rate. We highlight three: Meta-analysis, improved statistical practices, and machine learning.

For meta-analysis to succeed, the statistical effects from a large number of disparate studies must be collected, normalized, and reported in standardized ways¹⁰¹. The variability in analysis and reporting practices across the cognitive neuroscience literature can make meta-analysis challenging. As a result, the creation and curation of large-scale brain imaging databases has been essential for the growth of meta-analysis as an inferential tool. One of the oldest such systems devoted to supporting meta-analytic datasets and software is the BrainMap project (www.brainmap.org) project⁵. As of late fall 2016, BrainMap consisted of data from more than 100,000 individual participants from nearly 4,000 papers. The BrainMap data and meta-analysis tools have been used and cited more than 600 times since 1992, with more than 125 citations in 2016 alone.

Neurosynth (neurosynth.org)²⁶ takes an alternative approach to meta-analysis in which the raw data are i) activation (x,y,z) coordinates mined from the text of imaging papers published in HTML on the web, combined with ii) word frequencies from the same papers. In this way, Neurosynth aims to automate and thereby standardize and accelerate the process of meta-analysis. By combining information about activation coordinates with term frequencies derived from calculating distributional statistics from the published articles, Neurosynth enables the analyst to interactively determine the extent of evidence for a relationship between a specific term of interest and a set of brain coordinates. For example, the analyst could visualize either the probability of a given voxel's activation given the existence of a specific term in the system's database of papers or the probability of a target term appearing in papers that report a particular voxel as active. The system allows users to view 3D maps of the conditional probabilities online, to download the maps for further analysis, and to create customized sets of searches. As of

early 2017, Neurosynth contained data from more than 11,000 imaging studies, and it provides users with downloadable interactive meta-analyses from more than 3,000 terms. The system has been cited almost 600 times, with more than 180 citations in 2016 alone. Of course, Neurosynth only supports meta-analysis on a subset of the *published* literature—older papers that were not published in easily parsable HTML formats and unpublished findings fall outside its scope.

Beyond meta-analysis, cognitive neuroscience research continues to push for new statistical procedures and the wider adoption of long-standing, but more robust ones. For space reasons we do not elaborate extensively here, but among the issues under active discussion are the appropriate handling of main effects and interaction tests ¹⁰², the applicability of linear mixed effects modeling techniques ¹⁰³, and the ongoing need to guard against the risks of false positive results ^{3, 104} even when using well-established and vetted and widely used analysis software ¹². Still others suggest that the standard practice of treating stimulus effects as fixed, not random, may undermine the generalizability of findings across studies ¹⁰⁵. An emergent theme is the ongoing need for vigorous and rigorous methodological reevaluation combined with a commitment to more open software publication practices. In the recent case of Eklund and colleagues ¹², the discovery of an error in the algorithm for controlling for cluster-wise fMRI activation effects quickly led to changes in the widely-used AFNI package ¹⁰⁶⁻¹⁰⁷. The episode highlights the corrective, collaborative nature of open-source software development, while also underscoring the uncomfortable reality that, at present, very few people who use open-source software packages actually bother to read the underlying code (the AFNI bug had previously gone undetected for many years).

In many other areas of social and computational science, progress has been facilitated by borrowing ideas and techniques from the field of *machine learning*. Philosophically, machine learning researchers tend to emphasize their ability to quantitatively predict key outcomes and pay less attention to traditional forms of scientific explanation ¹⁰⁸. This philosophy has led to the rapid proliferation of thousands of predictive modeling techniques—a number of which (e.g., support vector machines) are deployed regularly in cognitive neuroscience. Many fMRI studies are now framed as predictive problems in classification or regression, where the goal is to build a model that successfully discovers a mapping between a set of predictor variables and a set of discrete (in the case of classification) or continuous (in the case of regression) outcomes. For example, the distributed activation pattern of a large number of voxels in an fMRI data set can be used to predict successful vs. unsuccessful attempts to recognize a stimulus. The resultant classifier can then be used to predict outcomes “out-of-sample” (i.e., in new data sets) and potentially also to aid in the interpretation of which voxels were likely to have played a role in processing relevant information. This gives rise to applications such as the ‘mind-reading’ of representations present in visual cortex during movie viewing ¹⁰⁹⁻¹¹⁰ or revealing participants’ semantic maps activated during narrative comprehension ¹¹¹.

Machine learning and related ‘big data’ techniques have provided entirely new approaches to analysis. They allow researchers to capitalize on neural information patterns that may be too subtle or complex to be easily discovered using more conventional summary statistic approaches (e.g. brain activation maps or ERPs) ¹¹². Of course, prediction-oriented approaches are not a panacea for standard concerns about the seeming ease with which researchers can fool themselves and unwittingly generate false or exaggerated findings. In the context of machine learning, the term ‘overfitting’ is used to describe a case in which the predictions of an

analysis have been inadvertently contaminated by noise in the data that were used to develop the analysis. As a consequence of overfitting, favorable results obtained when analyzing the same dataset that was used to develop and calibrate the analysis will not be obtained when examining other, equivalent datasets. Some researchers who use machine learning take the notion of ‘overfitting’ more seriously than others ¹¹³. The cautious deploy methods such as cross-validation (i.e., training and testing a model on independent subsets of the data) that should in principle guard against overfitting. However, machine learning pipelines allow for considerably more analytical flexibility than conventional analyses. Researchers often have a choice between literally hundreds of different estimation approaches, each of which may have its own free parameters that require tuning to perform optimally. Whereas there is widespread awareness of the need to cross-validate results once a model has been selected, there is much less recognition that overfitting can still occur through the optimization of an analysis (for further discussion, see ¹¹⁴⁻¹¹⁵). Thus, as our field increasingly adopts machine learning techniques, it will be important to borrow established best practices from fields that have been using similar big data approaches for longer periods of time ^{112,115}.

The Future

As Van Horn and Gazzaniga ⁵⁶ observed, “the reality remains that very little of the neuroimaging data gathered each day in the field have been made available to those who could help provide much needed understanding.” While we agree that this assessment still holds, we see other evidence that points toward a very different future. There is increasing recognition that greater openness and transparency, reflected in data, materials, and code sharing offers individual investigators and the field as a whole far more benefits than risks ^{28, 58, 108, 116}. While significant challenges remain in developing technology and workflow practices that make open and

transparent workflows easy to generate and data readily shareable, that progress is being made. It is increasingly clear that there are substantial scientific rewards in analyzing or re-analyzing large-scale shared data sets beyond improving statistical power. Accordingly, while we take seriously the concerns many have raised recently about the methods, results, and inferential reproducibility of our field, we encourage our colleagues to embrace the newly emerging open science practices with an optimistic mindset^{56, 117}, as there is so much more to gain than to lose.

At the same time, it is essential that the field identify barriers that stand in the way of a more open, transparent, and reproducible neuroscience of cognition. One clear gap is the difficulty of capturing and reporting reproducible information about tasks, displays, and analysis procedures, although new data and materials repositories like OSF and Databrary, emerging data standards (BIDS; BigEEG), and pipelines (EEGLAB; nipy) can play constructive roles. Another concerns the need to forge community consensus around a set of principles about the culture in which cognitive neuroscience research is carried out—how to seek permission to share, when data and materials should be shared, how to measure and report individual scholarly contributions to large-scale studies, how to weigh the impact of analyses conducted on secondary data relative to the collection of new data, and how to ensure that the transition to more open science practices doesn't unduly harm the careers of the next generation of researchers. Proscribing answers to these questions goes beyond our scope, but we urge continued dialogue focused on achieving community consensus.

A vital question for which there remains no satisfying answer is what entity will pay for the curation, support, maintenance and long-term storage of cognitive neuroscience data and

materials. Data repositories, both past and present, have been funded either by short-term (3-5 year duration) NIH or NSF research grants or private foundation funders. Thus, despite increasingly strong encouragements from granting agencies to share data and materials or even mandates to do so, there is as yet no long-term commitment from the agencies for funding devoted to long-term data preservation. Data curation, storage, preservation, is not inexpensive, and the problem of how to sustain research infrastructure that benefits the entire research community will neither solve itself nor go away. Nevertheless, based on the success of other fields like astronomy, high energy physics, and the geosciences, we think that a strong case can be made for enduring federal and private donor support for research infrastructure that empowers cognitive neuroscientists to openly share data, materials, and methods.

Fundamentally, we think that investments in the future of cognitive neuroscience infrastructure will generate big payoffs. Fostering the widespread adoption of open, transparent, and reproducible research practices coupled with innovations in technology that enable the large-scale analysis of our particular store of 'big data' will accelerate the discovery of generalizable, robust, and meaningful findings about the nature and origins¹¹⁸ of human cognition.

- 1 Bennett, C. M., & Miller, M. B. (2010). How reliable are the results from functional magnetic resonance imaging? *Annals of the New York Academy of Sciences*, 1191, 133–155. <http://doi.org/10.1111/j.1749-6632.2010.05446.x>
- 2 Carp, J. (2012b). The secret lives of experiments: Methods reporting in the fMRI literature. *NeuroImage*, 63(1), 289–300. <https://doi.org/10.1016/j.neuroimage.2012.07.004>
- 3 Vul, E., Harris, C., Winkielman, P., & Pashler, H. (2009). Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspectives on Psychological Science*, 4(3), 274–290. <http://doi.org/10.1111/j.1745-6924.2009.01125.x>
- 4 Biswal, B. B., Mennes, M., Zuo, X.-N., Gohel, S., Kelly, C., Smith, S. M., ... Milham, M. P. (2010). Toward discovery science of human brain function. *Proceedings of the National Academy of Sciences*, 107(10), 4734–4739. <https://doi.org/10.1073/pnas.0911855107>
- 5 Fox, P. T., Mikiten, S., Davis, G., & Lancaster, J. L. (1994). BrainMap: A database of human functional brain mapping. *Functional Neuroimaging: Technical Foundations*, 95-105. Chicago.
- 6 Van Horn, J. D., & Gazzaniga, M. S. (2002). Databasing fMRI studies — towards a “discovery science” of brain function. *Nature Reviews Neuroscience*, 3(4), 314–318. <http://doi.org/10.1038/nrn788>
- 7 Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), <http://doi.org/10.1126/science.aac4716>
- 8 Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365–376. <http://doi.org/10.1038/nrn3475>

- 9 David, S. P., Ware, J. J., Chu, I. M., Loftus, P. D., Fusar-Poli, P., Radua, J., ... Ioannidis, J. P. A. (2013). Potential reporting bias in fMRI studies of the brain. *PLoS ONE*, 8(7).
<http://doi.org/10.1371/journal.pone.0070104>
- 10 Ioannidis, J. P. A., Munafò, M. R., Fusar-Poli, P., Nosek, B. A., & David, S. P. (2014). Publication and other reporting biases in cognitive sciences: detection, prevalence, and prevention. *Trends in Cognitive Sciences*, 18(5), 235–241.
<http://doi.org/10.1016/j.tics.2014.02.010>
- 11 Szucs, D., & Ioannidis, J. P. (2016). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *bioRxiv*, 071530.
<https://doi.org/10.1101/071530>
- 12 Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences*, 113(28), 7900–7905. <https://doi.org/10.1073/pnas.1602413113>
- 13 Dolgin, E. (2010). This is your brain online: The functional connectomes project. *Nature Medicine*, 16(4), 351–351. <https://doi.org/10.1038/nm0410-351b>
- 14 Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., & Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80, 62–79.
<https://doi.org/10.1016/j.neuroimage.2013.05.041>
- 15 Adolph, K. E., Gilmore, R. O., Freeman, C., Sanderson, P., & Millman, D. (2012). Toward open behavioral science. *Psychological Inquiry*, 23(3), 244–247.
<https://doi.org/10.1080/1047840X.2012.705133>
- 16 Nosek, B. A., & Bar-Anan, Y. (2012). Scientific utopia: I. Opening scientific communication. *Psychological Inquiry*, 23(3), 217–243.
<https://doi.org/10.1080/1047840X.2012.692215>

- 17 Pashler, H., & Harris, C. R. (2012). Is the replicability crisis overblown? Three arguments examined. *Perspectives on Psychological Science*, 7(6), 531–536.
<https://doi.org/10.1177/1745691612463401>
- 18 Poldrack, R. A., & Poline, J.-B. (2015). The publication and reproducibility challenges of shared data. *Trends in Cognitive Sciences*, 19(2), 59–61.
<http://doi.org/10.1016/j.tics.2014.11.008>
- 19 Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Sert, N. P. du, ... Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, 1, 0021. <https://doi.org/10.1038/s41562-016-0021>
- 20 Goodman, S. N., Fanelli, D., & Ioannidis, J. P. A. (2016). What does research reproducibility mean? *Science Translational Medicine*, 8(341), 3
- 21 Nosek, B. A., Alter, G., Banks, G. C., Borsboom, D., Bowman, S. D., Breckler, S. J., ... Yarkoni, T. (2015). Promoting an open research culture. *Science*, 348(6242), 1422–1425.
<https://doi.org/10.1126/science.aab2374>
- 22 Patil, P., Peng, R. D., & Leek, J. T. (2016). What should researchers expect when they replicate studies? A statistical view of replicability in psychological science. *Perspectives on Psychological Science*, 11(4), 539–544. <https://doi.org/10.1177/1745691616646366>
- 23 Gilbert, D. T., King, G., Pettigrew, S., & Wilson, T. D. (2016). Comment on “Estimating the reproducibility of psychological science.” *Science*, 351(6277), 1037–1037.
<http://doi.org/10.1126/science.aad7243>
- 24 Etz, A., & Vandekerckhove, J. (2016). A Bayesian perspective on the Reproducibility Project: Psychology. *PLoS ONE*, 11(2), e0149794.
<https://doi.org/10.1371/journal.pone.0149794>

- 25 Yarkoni, T., Poldrack, R. A., Van Essen, D. C., & Wager, T. D. (2010). Cognitive neuroscience 2.0: Building a cumulative science of human brain function. *Trends in Cognitive Science*, 14(11), 489–496. <https://doi.org/10.1016/j.tics.2010.08.004>
- 26 Yarkoni, T., Poldrack, R. A., Nichols, T. E., Essen, D. C. V., & Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, 8(8), 665–670. <http://doi.org/10.1038/nmeth.1635>
- 27 Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific utopia II. Restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, 7(6), 615–631. <https://doi.org/10.1177/1745691612459058>
- 28 Poldrack, R. A., & Gorgolewski, K. J. (2014). Making big data open: Data sharing in neuroimaging. *Nature Neuroscience*, 17(11), 1510–1517. <http://doi.org/10.1038/nn.3818>
- 29 Poline, J.-B., Breeze, J. L., Ghosh, S. S., Gorgolewski, K., Halchenko, Y. O., Hanke, M., ... Kennedy, D. N. (2012). Data sharing in neuroimaging research. *Frontiers in Neuroinformatics*, 6(9). <https://doi.org/10.3389/fninf.2012.00009>
- 30 Wicherts, J. M., Bakker, M., & Molenaar, D. (2011). Willingness to share research data is related to the strength of the evidence and the quality of reporting of statistical results. *PLoS ONE*, 6(11), e26828. <https://doi.org/10.1371/journal.pone.0026828>
- 31 Luck, S. J. (2014). *An Introduction to the Event-Related Potential Technique*. MIT Press.
- 32 Delorme, A. & Makeig, S. (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics, *Journal of Neuroscience Methods*, 134, :9-21
- 33 Lopez-Calderon, J., & Luck, S. J. (2014). ERPLAB: An open-source toolbox for the analysis of event-related potentials. *Frontiers in Human Neuroscience*, 8(213). doi: 10.3389/fnhum.2014.00213

- 34 Poldrack, R. A., Barch, D. M., Mitchell, J., Wager, T., Wagner, A. D., Devlin, J. T., ... Milham, M. (2013). Toward open sharing of task-based fMRI data: the OpenfMRI project. *Frontiers in Neuroinformatics*, 7(12). <https://doi.org/10.3389/fninf.2013.00012>
- 35 practiCal fMRI: the nuts & bolts. (2013, January). A checklist for fMRI acquisition methods reporting in the literature. Retrieved from <https://practicalfmri.blogspot.com/2013/01/a-checklist-for-fmri-acquisition.html>
- 36 Gold, S., Christian, B., Arndt, S., Zeien, G., Cizadlo, T., Johnson, D. L., ... Andreasen, N. C. (1998). Functional MRI statistical software packages: A comparative analysis. *Human Brain Mapping*, 6(2), 73–84. [https://doi.org/10.1002/\(SICI\)1097-0193\(1998\)6:2<73::AID-HBM1>3.0.CO;2-H](https://doi.org/10.1002/(SICI)1097-0193(1998)6:2<73::AID-HBM1>3.0.CO;2-H) 41ps12–341ps12. <https://doi.org/10.1126/scitranslmed.aaf5027>
- 37 Morgan, V. L., Dawant, B. M., Li, Y., & Pickens, D. R. (2007). Comparison of fMRI statistical software packages and strategies for analysis of images containing random and stimulus-correlated motion. *Computerized Medical Imaging and Graphics*, 31(6), 436–446. <https://doi.org/10.1016/j.compmedimag.2007.04.002>
- 38 Strother, S. C., Anderson, J., Hansen, L. K., Kjems, U., Kustra, R., Sidtis, J., ... Rottenberg, D. (2002). The quantitative evaluation of functional neuroimaging experiments: The NPAIRS data analysis framework. *NeuroImage*, 15(4), 747–771. <https://doi.org/10.1006/nimg.2001.1034>
- 39 Strother, S., La Conte, S., Kai Hansen, L., Anderson, J., Zhang, J., Pulapura, S., & Rottenberg, D. (2004). Optimizing the fMRI data-processing pipeline using prediction and reproducibility performance metrics: I. A preliminary group analysis. *NeuroImage*, 23, Supplement 1, S196–S207. <https://doi.org/10.1016/j.neuroimage.2004.07.022>

- 40 Carp, J. (2012a). On the plurality of (methodological) worlds: estimating the analytic flexibility of fMRI experiments. *Brain Imaging Methods*, 6(149).
<https://doi.org/10.3389/fnins.2012.00149>
- 41 Glatard, T., Lewis, L. B., Ferreira da Silva, R., Adalat, R., Beck, N., Lepage, C., ... Evans, A. C. (2015). Reproducibility of neuroimaging analyses across operating systems. *Frontiers in Neuroinformatics*, 9. <https://doi.org/10.3389/fninf.2015.00012>
- 42 Turner, J. A., & Laird, A. R. (2011). The cognitive paradigm ontology: Design and application. *Neuroinformatics*, 10(1), 57–66. <https://doi.org/10.1007/s12021-011-9126-x>
- 43 Poldrack, R. A., Kittur, A., Kalar, D., Miller, E., Seppa, C., Gil, Y., ... Bilder, R. M. (2011). The cognitive atlas: toward a knowledge foundation for cognitive neuroscience. *Frontiers in Neuroinformatics*, 5(17). <https://doi.org/10.3389/fninf.2011.00017>
- 44 Gershon, R. C., Cella, D., Fox, N. A., Havlik, R. J., Hendrie, H. C., & Wagster, M. V. (2010). Assessment of neurological and behavioural function: the NIH Toolbox. *The Lancet Neurology*, 9(2), 138–139. [https://doi.org/10.1016/S1474-4422\(09\)70335-7](https://doi.org/10.1016/S1474-4422(09)70335-7)
- 45 Hall, D., Huerta, M. F., McAuliffe, M. J., & Farber, G. K. (2012). Sharing heterogeneous data: The National Database for Autism Research. *Neuroinformatics*, 10(4), 331–339.
<https://doi.org/10.1007/s12021-012-9151-4>
- 46 Casey, B. J., Cohen, J. D., O'Craven, K., Davidson, R. J., Irwin, W., Nelson, C. A., ... Turski, P. A. (1998). Reproducibility of fMRI results across four institutions using a spatial working memory task. *NeuroImage*, 8(3), 249–261. <https://doi.org/10.1006/nimg.1998.0360>
- 47 Friedman, L., Stern, H., Brown, G. G., Mathalon, D. H., Turner, J., Glover, G. H., ... Potkin, S. G. (2008). Test–retest and between-site reliability in a multicenter fMRI study. *Human Brain Mapping*, 29(8), 958–972. <https://doi.org/10.1002/hbm.20440>

- 48 Friedman, L., & Glover, G. H. (2006). Report on a multicenter fMRI quality assurance protocol. *Journal of Magnetic Resonance Imaging*, 23(6), 827–839.
<https://doi.org/10.1002/jmri.20583>
- 49 Brown, G. G., Mathalon, D. H., Stern, H., Ford, J., Mueller, B., Greve, D. N., ... Potkin, S. G. (2011). Multisite reliability of cognitive BOLD data. *NeuroImage*, 54(3), 2163–2175.
<https://doi.org/10.1016/j.neuroimage.2010.09.076>
- 50 Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. *NeuroImage*, 5(1), 49–62. <https://doi.org/10.1006/nimg.1996.0247>
- 51 Dosenbach, N. U. F., Visscher, K. M., Palmer, E. D., Miezin, F. M., Wenger, K. K., Kang, H. C., ... Petersen, S. E. (2006). A core system for the implementation of task sets. *Neuron*, 50(5), 799–812. <https://doi.org/10.1016/j.neuron.2006.04.031>
- 52 Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, 14(4), 172–179.
<https://doi.org/10.1016/j.tics.2010.01.004>
- 53 Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., ... Beckmann, C. F. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106(31), 13040–13045.
<https://doi.org/10.1073/pnas.0905267106>
- 54 Vigneau, M., Beaucoisin, V., Hervé, P. Y., Duffau, H., Crivello, F., Houdé, O., ... Tzourio-Mazoyer, N. (2006). Meta-analyzing left hemisphere language areas: Phonology, semantics, and sentence processing. *NeuroImage*, 30(4), 1414–1432.
<https://doi.org/10.1016/j.neuroimage.2005.11.002>

- 55 Yarkoni, T., & Braver, T. S. (2010). Cognitive neuroscience approaches to individual differences in working memory and executive control: Conceptual and methodological issues. In A. Gruszka, G. Matthews, & B. Szymura (Eds.), *Handbook of Individual Differences in Cognition*, 87–107. Springer New York. Retrieved from http://link.springer.com/chapter/10.1007/978-1-4419-1210-7_6
- 56 Van Horn, J. D., & Gazzaniga, M. S. (2013). Why share data? Lessons learned from the fMRIDC. *NeuroImage*, 82, 677–682. <https://doi.org/10.1016/j.neuroimage.2012.11.010>
- 57 Van Horn, J.D., Grethe, J.S., Kostelec, P., Woodward, J.B., Aslam, J.A., Rus, D., Rockmore, D., Gazzaniga, M.S. (2001). The Functional Magnetic Resonance Imaging Data Center (fMRIDC): The challenges and rewards of large-scale databasing of neuroimaging studies. *Philosophical Transactions of the Royal Society of London: B Biological Sciences*, 356, 1323–1339.
- 58 Mennes, M., Biswal, B., Castellanos, F. X., & Milham, M. P. (2013). Making data sharing work: The FCP/INDI experience. *NeuroImage*, 82, 683–691. <http://doi.org/10.1016/j.neuroimage.2012.10.064>
- 59 Miller, M. B., Donovan, C.-L., Van Horn, J. D., German, E., Sokol-Hessner, P., & Wolford, G. L. (2009). Unique and persistent individual patterns of brain activity across different memory retrieval tasks. *NeuroImage*, 48(3), 625–635. <https://doi.org/10.1016/j.neuroimage.2009.06.033>
- 60 Costafreda, S. G., Brammer, M. J., Vêncio, R. Z. N., Mourão, M. L., Portela, L. A. P., de Castro, C. C., ... Amaro, E. (2007). Multisite fMRI reproducibility of a motor task using identical MR systems. *Journal of Magnetic Resonance Imaging*, 26(4), 1122–1126. <https://doi.org/10.1002/jmri.21118>

- 61 Duncan, K. J., Pattamadilok, C., Knierim, I., & Devlin, J. T. (2009). Consistency and variability in functional localisers. *NeuroImage*, 46(4), 1018–1026.
<https://doi.org/10.1016/j.neuroimage.2009.03.014>
- 62 Poldrack, R. A., Baker, C. I., Durnez, J., Gorgolewski, K. J., Matthews, P. M., Munafò, M. R., ... Yarkoni, T. (2017). Scanning the horizon: towards transparent and reproducible neuroimaging research. *Nature Reviews Neuroscience*, advance online publication.
<https://doi.org/10.1038/nrn.2016.167>
- 63 Thirion, B., Pinel, P., Mériaux, S., Roche, A., Dehaene, S., & Poline, J.-B. (2007). Analysis of a large fMRI cohort: Statistical and methodological issues for group analyses. *NeuroImage*, 35(1), 105–120. <https://doi.org/10.1016/j.neuroimage.2006.11.054>
- 64 Yarkoni, T. (2009). Big Correlations in Little Studies: Inflated fMRI Correlations Reflect Low Statistical Power—Commentary on Vul et al. (2009). *Perspectives on Psychological Science*, 4(3), 294–298. <https://doi.org/10.1111/j.1745-6924.2009.01127.x>
- 65 Kriegeskorte, N., Lindquist, M. A., Nichols, T. E., Poldrack, R. A., & Vul, E. (2010). Everything You Never Wanted to Know about Circular Analysis, but Were Afraid to Ask. *Journal of Cerebral Blood Flow & Metabolism*, 30(9), 1551–1557. <https://doi.org/10.1038/jcbfm.2010.86>
- 66 Bennett, C., Miller, M., & Wolford, G. (2009). Neural correlates of interspecies perspective taking in the post-mortem Atlantic Salmon: an argument for multiple comparisons correction. *NeuroImage*, 47, Supplement 1, S125.
[https://doi.org/10.1016/S1053-8119\(09\)71202-9](https://doi.org/10.1016/S1053-8119(09)71202-9)
- 67 Bennett, C. M., Wolford, G. L., & Miller, M. B. (2009). The principled control of false positives in neuroimaging. *Social Cognitive and Affective Neuroscience*, 4(4), 417–422.
<https://doi.org/10.1093/scan/nsp053>

- 68 Woo, C.-W., Krishnan, A., & Wager, T. D. (2014). Cluster-extent based thresholding in fMRI analyses: Pitfalls and recommendations. *NeuroImage*, 91, 412–419.
<https://doi.org/10.1016/j.neuroimage.2013.12.058>
- 60 Smith, S. M., & Nichols, T. E. (2009). Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *NeuroImage*, 44(1), 83–98. <https://doi.org/10.1016/j.neuroimage.2008.03.061>
- 70 Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), e124. <https://doi.org/10.1371/journal.pmed.0020124>
- 71 Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(7197), 869–878. <https://doi.org/10.1038/nature06976>
- 72 Cole, D. M., Smith, S. M., & Beckmann, C. F. (2010). Advances and pitfalls in the analysis and interpretation of resting-state FMRI data. *Frontiers in Systems Neuroscience*, 4. <https://doi.org/10.3389/fnsys.2010.00008>
- 73 Poldrack, R. A. (2010). Subtraction and beyond: The logic of experimental designs for neuroimaging. In S. J. Hanson & M. Bunzl (Eds.), *Foundational Issues in Human Brain Mapping*. MIT Press.
- 74 Etzel, J. A., Zacks, J. M., & Braver, T. S. (2013). Searchlight analysis: Promise, pitfalls, and potential. *NeuroImage*, 78, 261–269. <https://doi.org/10.1016/j.neuroimage.2013.03.041>
- 75 Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Sciences*, 10(2), 59–63. <https://doi.org/10.1016/j.tics.2005.12.004>
- 76 Peng, R. D. (2011). Reproducible research in computational science. *Science*, 334(6060), 1226–1227. <https://doi.org/10.1126/science.1213847>
- 77 Stodden, V. (2012). Reproducible research: Tools and strategies for scientific computing. *Computing in Science & Engineering*, 14(4), 11–12. <https://doi.org/10.1109/MCSE.2012.82>

- 78 Sandve, G. K., Nekrutenko, A., Taylor, J., & Hovig, E. (2013). Ten Simple Rules for Reproducible Computational Research. *PLoS Comput Biol*, 9(10), e1003285.
<https://doi.org/10.1371/journal.pcbi.1003285>
- 79 Gorgolewski, K. J., Auer, T., Calhoun, V. D., Craddock, R. C., Das, S., Duff, E. P., ... Poldrack, R. A. (2016). The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data*, 3, 160044.
<https://doi.org/10.1038/sdata.2016.44>
- 80 Gorgolewski, K. J., Varoquaux, G., Rivera, G., Schwarz, Y., Ghosh, S. S., Maumet, C., ... Margulies, D. S. (2015). NeuroVault.org: a web-based repository for collecting and sharing unthresholded statistical maps of the human brain. *Frontiers in Neuroinformatics*, 9.
<https://doi.org/10.3389/fninf.2015.00008>
- 81 Gorgolewski, K. J., Alfaro-Almagro, F., Auer, T., Bellec, P., Capota, M., Chakravarty, M., ... Poldrack, R. (2016). BIDS Apps: Improving ease of use, accessibility and reproducibility of neuroimaging data analysis methods. *bioRxiv*, 079145. <https://doi.org/10.1101/079145>
- 82 Gilmore, R.O., & Adolph, K.E. (in press). Open Sharing of Research Video -- Breaking down the boundaries of the research team, in Hall, Kara, Croyle, R., & Vogel, A. (Eds.), *Advancing Social and Behavioral Health Research through Cross-Disciplinary Team Science: Principles for Success*. Springer.
- 83 Jann, K., Gee, D. G., Kilroy, E., Schwab, S., Smith, R. X., Cannon, T. D., & Wang, D. J. (2015). Functional connectivity in BOLD and CBF data: Similarity and reliability of resting brain networks. *NeuroImage*, 106, 111–122. <https://doi.org/10.1016/j.neuroimage.2014.11.028>
- 84 Jovicich, J., Marizzoni, M., Bosch, B., Bartrés-Faz, D., Arnold, J., Benninghoff, J., ... Frisoni, G. B. (2014). Multisite longitudinal reliability of tract-based spatial statistics in diffusion

tensor imaging of healthy elderly subjects. *NeuroImage*, 101, 390–403.

<https://doi.org/10.1016/j.neuroimage.2014.06.075>

85 Koolschijn, P. C. M. P., Schel, M. A., Rooij, M. de, Rombouts, S. A. R. B., & Crone, E. A. (2011). A three-year longitudinal functional magnetic resonance imaging study of performance monitoring and test-retest reliability from childhood to early adulthood. *Journal of Neuroscience*, 31(11), 4204–4212. <https://doi.org/10.1523/JNEUROSCI.6415-10.2011>

86 Liao, X.-H., Xia, M.-R., Xu, T., Dai, Z.-J., Cao, X.-Y., Niu, H.-J., ... He, Y. (2013). Functional brain hubs and their test–retest reliability: A multiband resting-state functional MRI study. *NeuroImage*, 83, 969–982. <https://doi.org/10.1016/j.neuroimage.2013.07.058>

87 Madhyastha, T., Mérillat, S., Hirsiger, S., Bezzola, L., Liem, F., Grabowski, T., & Jäncke, L. (2014). Longitudinal reliability of tract-based spatial statistics in diffusion tensor imaging. *Human Brain Mapping*, 35(9), 4544–4555. <https://doi.org/10.1002/hbm.22493>

88 Marchitelli, R., Minati, L., Marizzoni, M., Bosch, B., Bartrés-Faz, D., Müller, B. W., ... Jovicich, J. (2016). Test-retest reliability of the default mode network in a multi-centric fMRI study of healthy elderly: Effects of data-driven physiological noise correction techniques. *Human Brain Mapping*, 37(6), 2114–2132. <https://doi.org/10.1002/hbm.23157>

89 Choe, A. S., Jones, C. K., Joel, S. E., Muschelli, J., Belegu, V., Caffo, B. S., ... Pekar, J. J. (2015). Reproducibility and temporal structure in weekly resting-state fMRI over a period of 3.5 years. *PLoS ONE*, 10(10), e0140134. <http://doi.org/10.1371/journal.pone.0140134>

90 Poldrack, R. A., Laumann, T. O., Koyejo, O., Gregory, B., Hover, A., Chen, M.-Y., ... Mumford, J. A. (2015). Long-term neural and physiological phenotyping of a single human. *Nature Communications*, 6, 8885. <https://doi.org/10.1038/ncomms9885>

91 Dubois, J., & Adolphs, R. (2016). Building a science of individual differences from fMRI. *Trends in Cognitive Sciences*, 20(6), 425–443. <https://doi.org/10.1016/j.tics.2016.03.014>

- 92 Boekel, W., Wagenmakers, E.-J., Belay, L., Verhagen, J., Brown, S., & Forstmann, B. U. (2015). A purely confirmatory replication study of structural brain-behavior correlations. *Cortex*, 66, 115–133. <https://doi.org/10.1016/j.cortex.2014.11.019>
- 93 Muhlert, N., & Ridgway, G. R. (2016). Failed replications, contributing factors and careful interpretations: Commentary on Boekel et al., 2015. *Cortex*, 74, 338–342. <https://doi.org/10.1016/j.cortex.2015.02.019>
- 94 Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., ... Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615), 171–178. <https://doi.org/10.1038/nature18933>
- 95 Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., ... Jenkinson, M. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80, 105–124. <https://doi.org/10.1016/j.neuroimage.2013.04.127>
- 96 Zuo, X.-N., Anderson, J. S., Bellec, P., Birn, R. M., Biswal, B. B., Blautzik, J., ... Milham, M. P. (2014). An open science resource for establishing reliability and reproducibility in functional connectomics. *Scientific Data*, 1, 140049. <https://doi.org/10.1038/sdata.2014.49>
- 97 Nichols, T. E., Das, S., Eickhoff, S. B., Evans, A. C., Glatard, T., Hanke, M., ... Yeo, B. T. T. (2016). Best practices in data analysis and sharing in neuroimaging using MRI. *bioRxiv*, 054262. <https://doi.org/10.1101/054262>
- 98 Epskamp, S. & Nuijten, M. B. (2016). statcheck: Extract statistics from articles and recompute p values. Retrieved from <http://CRAN.R-project.org/package=statcheck>.
- 99 Nuijten, M. B., Hartgerink, C. H. J., Assen, M. A. L. M. van, Epskamp, S., & Wicherts, J. M. (2015). The prevalence of statistical reporting errors in psychology (1985–2013). *Behavior Research Methods*, 48(4), 1–22. <https://doi.org/10.3758/s13428-015-0664-2>

- 100 Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve: a key to the file-drawer. *Journal of Experimental Psychology: General*, 143(2), 534.
- 101 Wager, T. D., Lindquist, M., & Kaplan, L. (2007). Meta-analysis of functional neuroimaging data: current and future directions. *Social, Cognitive and Affective Neuroscience*, 2(2), 150–158. <https://doi.org/10.1093/scan/nsm015>
- 102 Nieuwenhuis, S., Forstmann, B. U., & Wagenmakers, E.-J. (2011). Erroneous analyses of interactions in neuroscience: a problem of significance. *Nature Neuroscience*, 14(9), 1105–1107. <https://doi.org/10.1038/nn.2886>
- 103 Chen, G., Saad, Z. S., Britton, J. C., Pine, D. S., & Cox, R. W. (2013). Linear mixed-effects modeling approach to fMRI group analysis. *NeuroImage*, 73, 176–190. <https://doi.org/10.1016/j.neuroimage.2013.01.047>
- 104 Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. F., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nature Neuroscience*, 12(5), 535–540. <https://doi.org/10.1038/nn.2303>
- 105 Westfall, J., Nichols, T., & Yarkoni, T. (2016). Fixing the stimulus-as-fixed-effect fallacy in task fMRI. *bioRxiv*, 077131. <https://doi.org/10.1101/077131>
- 106 Cox, R. W., Reynolds, R. C., & Taylor, P. A. (2016). AFNI and clustering: False positive fates fedux. *bioRxiv*, 065862. <https://doi.org/10.1101/065862>
- 107 Eickhoff, S. B., Laird, A. R., Fox, P. M., Lancaster, J. L., & Fox, P. T. (2016). Implementation errors in the GingerALE Software: Description and recommendations. *Human Brain Mapping*, 38(1). <https://doi.org/10.1002/hbm.23342>
- 108 Breiman, L. (2001). Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3), 199–231. <https://doi.org/10.1214/ss/1009213726>

- 109 Kay, K. N., Naselaris, T., Prenger, R. J., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, 452(7185), 352–355.
<https://doi.org/10.1038/nature06713>
- 110 Nishimoto, S., Vu, A. T., Naselaris, T., Benjamini, Y., Yu, B., & Gallant, J. L. (2011). Reconstructing visual experiences from brain activity evoked by natural movies. *Current Biology*, 21(19), 1641–1646. <https://doi.org/10.1016/j.cub.2011.08.031>
- 111 Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E., & Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532(7600), 453–458. <https://doi.org/10.1038/nature17637>
- 112 Varoquaux, G., Raamana, P. R., Engemann, D., Hoyos-Idrobo, A., Schwartz, Y., & Thirion, B. (2017). Assessing and tuning brain decoders: Cross-validation, caveats, and guidelines. *NeuroImage*, 145, 166–179. <https://doi.org/10.1016/j.neuroimage.2016.10.038>
- 113 Yarkoni, T., & Westfall, J. (in press). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*.
- 114 Cawley, G. C., & Talbot, N. L. C. (2010). On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11(Jul), 2079–2107. <http://www.jmlr.org/papers/v11/cawley10a.html>
- 115 Skocik, M., Collins, J., Callahan-Flintoft, C., Bowman, H., & Wyble, B. (2016). I tried a bunch of things: The dangers of unexpected overfitting in classification. *bioRxiv*, 078816.
- 116 Bandrowski, A. E., & Martone, M. E. (2016). RRIDs: A simple step toward improving reproducibility through rigor and transparency of experimental methods. *Neuron*, 90(3), 434–436. <http://doi.org/10.1016/j.neuron.2016.04.030>
- 117 Ascoli, G. A. (2006). The ups and downs of neuroscience shares. *Neuroinformatics*, 4(3), 213–215. <https://doi.org/10.1385/NI:4:3:213>

118 Gilmore, R.O. (2016). From big data to deep insight in developmental science. *Wiley Interdisciplinary Reviews Cognitive Science*. doi: 10.1002/wcs.1379.

Author Note

ROG acknowledges support from NSF BCS-1147440, NSF BCS-1238599, and NICHD

U01-HD-076595. BW acknowledges support from NSF BCS-1331073