

DOI: 10.1017/psa.2025.14

This is a manuscript accepted for publication in *Philosophy of Science*.

This version may be subject to change during the production process.

Trading Evidence: The Role of Models in Interfield Unification

Daniel A. Weiskopf

Contact: Department of Philosophy, Georgia State University, PO Box 3994, Atlanta, GA 30302-3984. Email: dweiskopf@gsu.edu

Abstract: Scientific fields frequently need to exchange data to advance their own inquiries. Data unification is the process of stabilizing these forms of interfield data exchange. I present an account of the epistemic structure of data unification, drawing on case studies from model-based cognitive neuroscience (MBCN). MBCN is distinctive because it shows that modeling practices play an essential role in mediating these data exchanges. Models often serve as interfield evidential integrators, and models built for this purpose have their own representational and inferential functions. This form of data unification should be seen as autonomous from other forms, particularly explanatory unification.

Acknowledgments: Versions of this paper were presented at the Evidence and Explanations of Cognition conference hosted by the University of Kent, at the Cognitive Science Seminar of the UNAM Philosophy Department, and at the Deep South Philosophy and Neuroscience Workgroup. I am grateful to the participants for their helpful discussion, as well as to Zoe Drayson for comments on the draft. Work on this project was supported by a GSU Research Innovation and Scholarly Excellence grant in Trustworthy Artificial Intelligence for an Equitable Digital Society.

1. Forms of unification

Accounts of scientific unification differ in what they take the target of unification to be. Historically, unification has centered on abstract representations such as theories, models, and explanations more generally. This trend can be seen in classic work on intertheoretic reduction, explanatory unification, and mechanistic explanation, all of which are concerned largely with explicating the senses in which these epistemic products can themselves be unified, or play a central role in unifying fields.¹ Such unification takes many forms, including showing that one field's theories can be subsumed under another, that models from one field can be transferred to another, or that information from multiple lines of inquiry can contribute to building multilevel models that span the domains of several related fields.

Contemporary work on unification, however, has drawn attention to other ways that scientific communities can productively merge their activities. Fields frequently exchange not just theories and models but also instruments, tools, experimental systems, computational or analytic techniques, and so on. These are not products of inquiry, but methods and materials developed to support inquiry. Trading and adapting these materials promotes a distinctive kind of field unification, since the use of shared techniques and methods need not lead to unified theories and models. It is possible, indeed commonplace, to adopt a set of statistical methods without using them to support the same kinds of theories, or to repurpose instruments invented for one kind of inquiry to serve another. Methods, tools, and instruments are flexible with respect to the goals they can serve.

In addition to having their own explanations and methods, fields generate their own characteristic types of data. Data occupy a position between these two more widely explored forms of scientific unification insofar as they are generated by experimental systems and instruments, processed using particular analytic methods, and ultimately serve to support theories, models, and explanations. Just as explanations and methods may migrate across fields, so too can data. We therefore have a threefold way of classifying scientific unification in terms

¹ Here I will treat the social units of scientific organization as fields rather than disciplines, in roughly Darden's sense (Darden 2005; Darden and Maull 1977).

of explanations, methods, and data.² Understanding data unification is particularly important within massively interfield projects such as the cognitive sciences. Here the typical data are highly varied and may include grammaticality judgments and corpus statistics, response times and error rates, eye movement tracks and pupil dilation measurements, looking times, stimulus similarity ratings, verbal protocols, spike train recordings from single-cell electrophysiology, event-related potentials, single photon emission CT and functional MRI images, and so on. Investigators often want to use data and observations collected within one field to investigate questions posed by another. White matter tractography maps from neuroimaging may inform theorizing in cognitive psychology; cross-cultural ethnographic observations may inform developmental theorizing; receptor anatomy and psychophysical response times may inform speculations about neural evolution.

It is often unclear, however, how to use data from outside of one field to help settle questions within it. For that matter, it may not be obvious whether these data even can be used in this way. The evidential value of data—the hypotheses they may support and questions they may address—is not written on their face. Theorists of interfield unification and integration have recently begun to describe the structure of these data exchanges, giving rise to a number of accounts of how data from one field are turned into evidence for another (Leonelli 2016, 141–75; O'Malley 2013; O'Malley and Soyer 2012). Here I contribute to this project by arguing that models play a central role in enabling interfield data exchanges. Specifically, model-building practices can serve as sites for passing data across fields and thereby coordinating fields' epistemic activities.

I first introduce a conception of data as varying in their evidential value depending on the context of their use (Section 2). With this conception in hand, the process of model-based coordination becomes one of turning data from one field into something having evidential value for another. This shift in epistemic value is illustrated by looking at a set of cases at the nexus of experimental psychology and neuroscience, specifically in the emerging interfield practices of model-based cognitive neuroscience (Section 3). Analyzing these cases will highlight the complex representational and inferential structure of the models involved in evidential

² I owe the distinctions among methodological, data, and explanatory integration to O'Malley and Soyer's (2012) excellent and detailed discussion.

coordination (Section 4). This form of interfield data unification can be understood as a kind of practical rather than theoretical unification (Section 5). Finally, data unification can be productively distinguished from neighboring notions of explanatory unification (Section 6).

2. Sharing data and making evidence

To begin we will need a more precise conception of what data are. A prominent account owing to Woodward (1989, 394) holds that “data are what registers on a measurement or recording device in a form which is accessible to the human perceptual system, and to public inspection,” as well as the “records or reports” of such events (Woodward 2010, 792). On this view, data are marks, records, or traces of causal processes. The epistemology of data therefore centers whether they are reliable in this role, looking to the experimental setups, instruments, and intervention techniques that have produced them to winnow out possible sources of distortion, confounds, noise, bias, and artefacts. The assessment of new measurement techniques in neuroscience, for example, has often focused on assessing the methods by which data are generated (Bechtel Forthcoming) with the aim of permitting reliable discrimination among competing claims about the phenomena (Woodward 2000). Data on this picture are tightly linked with the specific and sometimes ill-understood causal setups that generate them. The more secure these connections, the more confidently we can take data to be informative about some aspect of the world.

Leonelli (2016, 77) has proposed a relational conception of data on which they are “any product of research activities, ranging from artifacts such as photographs to symbols such as letters or numbers, that is collected, stored, and disseminated *in order to be used as evidence for knowledge claims*.” She does not discard questions about the origins of data—in fact, keeping track of such historical facts is essential to creating data histories, which play a key role in determining the weight we should give data in different kinds of inquiry.³ The relational

³ As she notes, evidential value is “judged by scientists through an evaluation of the ways in which data have been collected and disseminated, including the instruments and materials employed to that effect” (2013, 3).

These evaluations are part of what allows data to convey their potential evidential value in contexts outside of the ones where they were generated. On the granular differences between Leonelli’s conception and Woodward’s, see her 2016, 84-8.

conception, however, emphasizes two central traits of data that go beyond backwards-looking assessment. First, they are portable or detachable from their circumstances of production. Data are by their nature things that can circulate within and across research communities. In circulating, they become available to be taken up for many purposes that can differ from those of their creators: they have the “ability to fit a variety of lines of inquiry” (80). Second, data depend for their usefulness on their reception or circumstances of application “because one cannot predict *all* the possible claims that data might be used as evidence for in the future; and also because one cannot predict whether data will actually be used as evidence for specific claims until it *happens*” (84). As Leonelli’s characterization implies, data are not automatically evidence. Something’s status as data depends on its potential to serve as evidence relative to an inquiry. Data become evidence for a field when there are reliable ways to use them to support, reject, constrain, or extend claims made within that field. The work of evidential exchange lies in this activity, particularly when the data are generated using tools and methods that come from separate fields.

The relational conception shifts our attention from the circumstances of data production to those of reception. It asks us to assess data in terms of the potential breadth of their relevance. Within the relational conception it is not always obvious whether some item is data or not, that is, whether it has potential evidential value. Countless things can be reliably measured, but those measurements may turn out to be worthless. Nor it is obvious how a putative piece of data’s evidential value can be exploited by using it to systematically do epistemic work. The problem faced in interfield inquiry is not always (or even primarily) in collecting new data but in establishing these two kinds of relevance relations and thereby showing how claims within a field can be potentially confirmed or undermined by data imported from outside of it.

The most promising aspects of these two conceptions of data have been synthesized by Bokulich and Parker (2021) in their “pragmatic-representational” view. They take data to be “records of the results of a process of inquiry that involves interacting with the world” (6) that are assessed according to how adequate they are for specific purposes.⁴ Adequacy-for-purpose is

⁴ Bokulich and Parker further argue that insofar as data can be taken to be about some aspect of the world that is involved in their production, they can be construed as representations. They add that while the

a multivalent property that depends on whether the data are actually usable for achieving particular situational ends, or whether they would be usable given the right set of informational, cognitive, or technological resources (11). A data set can therefore be evaluated differently by two investigators depending on how they value qualities such as accuracy or high resolution, or which sets of inferences they are intending to use it to make. Crucially, data can be dynamically repurposed, that is, used to answer an entirely new set of questions given the right sort of processing (19-20). Repurposing can involve using a data set to draw conclusions about an entirely different target than the one it originally concerned: in their example, researchers used Mars rover navigational data to measure regional gravitational changes and ultimately to estimate the density of the Martian crust at different locations. This sort of chained repurposing illustrates the mobile quality of data that is also central to Leonelli's relational account.

Discovering ways to repurpose data in order to answer new questions across fields is one way of making that data evidentially relevant. Here I will adopt the pragmatic-representational conception as a way of understanding interfield data unification in the cognitive sciences. The case studies below draw on and develop Leonelli's (2016, 91) insight that, very often, "[m]odeling is the process by which data are assigned evidential value."

3. Studies in model-based cognitive neuroscience

Model-based cognitive neuroscience (MBCN), according to its leading proponents and practitioners, comprises "several entirely new statistical modeling approaches developed through collaborations between mathematical psychologists and cognitive neuroscientists, collectively forming a new field" (Turner et al. 2017, 66) These approaches, mainly developed within the past 15 years, have been accompanied by the appearance of major textbooks, special journal issues, and edited volumes.

Putting aside for now the question of whether MBCN itself truly constitutes a new or emergently stabilizing field, I will treat it *pro tem* as a heterogeneous set of tools, analytic methods, experimental systems, and data sources that demonstrate all three forms of unification

representational value of data is constrained by the circumstances of their production, they need not represent *only* those aspects. See also Leonelli (2019, 17–19).

sketched above.⁵ From mathematical psychology it takes formal models of cognition such as the diffusion drift model (to be discussed shortly). From experimental psychology come protocols, methods, and intervention techniques for eliciting data about how human participants perform in various tasks that are hypothetically linked to the constructs described in these formal models. From cognitive neuroscience come technologies and experimental systems for measuring kinds of neural activity during the performance of these tasks. Finally, from statistics come a range of analytic methods and computational tools that are essential for using mathematical models and data sets to answer particular kinds of research questions.

MBCN represents, then, a highly active zone of epistemic trade.⁶ This activity is best appreciated by observing its practical implementations. The most lively domains of investigation in MBCN center on psychological capacities and phenomena for which there are existing mathematical models that have been built, trained, and tested in a range of behavioral experiments. These include models of reinforcement learning, categorization, and decision and choice. Here we will focus on studies of choice that draw on accumulator models of various types, best exemplified by (but not exclusive to) the diffusion decision model, sometimes also called the diffusion drift model (DDM).

First developed in a landmark paper by Roger Ratcliff (1978), the DDM in its elementary form is a model of two-alternative forced choice. Participants are shown a set of materials and given a prompt that requires making a binary decision such as “Are the dots in this image predominantly moving upwards or downwards?” or “Are these two images matched in their overall brightness or not?” The model represents such choices as resulting from a continuous (i.e., nondiscrete) process of evidence accumulation. Participants begin decision making at a starting point, which may be either neutral or biased towards one of the options. They are then presented with a changing stream of evidence that nudges them towards one option or the other. The process continues until the accumulated evidence reaches a decision threshold; crossing the threshold constitutes the act of choice and triggers a behavioral response.

⁵ MBCN not been extensively theorized by philosophers, but for some illuminating previous discussions see Irvine (2016) and Povich (2019).

⁶ See Section 4 for more on this notion and its relation to Peter Galison’s conception of trading zones.

The model in a streamlined form thus has four parameters interpretable as separate cognitive aspects of the decision process (Alexandrowicz 2020): (1) the location of the starting point between the two choice options; (2) the location of the decision boundaries for the options, which stand for how much evidence needs to be acquired to trigger a decision; (3) the drift rate that determines how much the choice point moves on average for each new piece of evidence acquired; and (4) the non-decision time, which collects all of the other sensory and motor factors that are extrinsic to the choice process itself. These parameters are usually assumed to vary on a trial-by-trial basis within participants. Setting these four parameters enables the DDM to be used to jointly predict participants' error rates and response times, as well as speed/accuracy tradeoffs and interactions with properties of the task.

Since its inception, the DDM has been applied across a wide range of conditions and experimental materials, including tasks such as recognition memory, perceptual decision, lexical and semantic decision, implicature detection, and consumer product choice (Forstmann et al. 2016; Ratcliff et al. 2016). It has been applied in aging, sleep-deprived, and psychiatric and neuroatypical populations (persons with aphasia, dyslexia, mood disorders, etc.) as well as to studies of individual differences (Ratcliff and Childers 2015).

Despite its popularity and successes, the model also faces challenges. One in particular, which Ratcliff (1978) calls the problem of model freedom, is endemic to cognitive models that draw on behavioral data. Such data are frequently not rich enough to constrain parameter values, meaning that many different parameter settings are compatible with similar behavioral predictions. Insufficient data also means that there will be limits on the precision of the model's predictions. If model parameters are intended to correspond to cognitive constructs, insufficient data also limits the usefulness of behaviorally-fitted models for describing the causal sources of those behaviors. Finally, limited data can make model selection challenging in cases where models with different structure make nearly identical behavioral predictions.

MBCN offers techniques to overcome these challenges. Neural data can provide greater constraints on parameter settings, more accurate neural and behavioral predictions, more precise understandings of the cognitive processes that the model represents, and potentially greater

resources for selecting among cognitive models. In what follows I outline some ways the toolkit of MBCN techniques contributes to answering these questions.⁷

3.1. Parameter estimation

One challenge in cognitive modeling is parameter estimation, which aims not only to produce the closest fit to how participants perform, but also to recover the best approximations to the values of the model parameters that generate that fit (so-called “ground truth”). Getting more accurate estimates of the roles of various constructs in cognition is a primary goal of this form of modeling, although approximately true parameter values may also make the model useful for other purposes, such as generating better predictions. Suppose one wants to answer a diagnostically important question such as whether longer response times in a specific population such as older adults is due to greater response caution or slowed motor execution. The literature on parameter estimation in the context of DDM has focused on questions such as how reliable empirical parameter estimates are (Lerche and Voss 2017), how model parameters should be understood and visualized (Alexandrowicz 2020), and what mathematical methods are optimal for estimating parameters (Alexandrowicz and Gula 2020; Arnold et al. 2015; Grasman et al. 2009). But models such as the DDM can address these questions only on the assumption that their tuned parameters track ground truths.

MBCN aims to overcome the parameter estimation problem by using neural data to improve accuracy. One technique for doing this is the joint modeling approach presented by Turner et al. (2016). Joint models are a type of Bayesian hierarchical model in which various submodels are linked to one another via hyperparameters. Formally, hyperparameters function to set the value of other model parameters. Within joint modeling, they are usually interpreted as encoding a hypothesis about the statistical connections among the parameters of the submodels. The particular form of the hyperparameters determines the kind and magnitude of this relationship. For this reason, hyperparameter choices are said to constitute linking propositions.⁸

⁷ For background on many of the techniques discussed here, particularly joint modeling in MBCN, see de Hollander et al. (2016), Turner et al. (2019) and Turner et al. (2013). For a review of hierarchical modeling in a cognitive context, see Shiffrin et al. (2008).

⁸ In MBCN, linking propositions are usually meant in Teller’s (1984, 1235) sense of the term: “a claim that a particular mapping occurs, or a particular mapping principle applies, between perceptual and physiological

So, for example, in the model to be discussed below, changes in BOLD values can be systematically related to both changes in drift rate and to EEG values, and therefore encode different potential ways of linking these quantities.

Turner et al.'s studies collected participants' response time data from a set of choice tasks meant to estimate each individual's temporal discount rate. Individuals carried out the tasks in one of several conditions: purely behavioral, behavioral with EEG recording only, with fMRI only, or with both EEG and fMRI. These data sets were then used to construct three submodels that feed into their final joint model. The cognitive submodel was a form of the Linear Ballistic Accumulator trained on the response time and error rate data.⁹ The LBA is another sequential sampler model like the DDM. While there are differences between LBA and DDM, they share a drift rate parameter that governs the cognitive process of evidence accumulation and so are for present purposes comparable.

The other two submodels included data from the two neurophysiological measures. A key issue in joint modeling is whether to include complete or reduced sets of neural recordings. Because previous work had established dorsomedial frontal cortex (dmFC) as significant for decision tasks, the EEG submodel was simplified to include only recordings from four electrodes located over dmFC. The fMRI model similarly included only BOLD activity measured from the same region. Specifically, the imaging data took the form of the single-trial β weights derived from fitting the generalized linear model to the fMRI scans. The two neural submodels (EEG and fMRI) thus contain considerably slimmed-down data models governed by just five parameters, where these selections were hypothesized to be relevant to the task based on prior knowledge.

The complete model enables us to pose questions about how two types of neural data relate to each other and to the cognitive parameters underlying behavior. The function of the

states.” As I have reconstructed things here, it is more accurate to think of these propositions as mapping not types of states, but forms of data. Linking propositions are assumptions about how data sources are related to one another in the context of a model.

⁹ Turner et al. call this a “behavioral” model, but since its parameters track cognitive constructs it should more appropriately be considered a cognitive model like the DDM. Accordingly, despite the potential for some confusion, I have renamed it here to emphasize the difference between thinking of the LBA as a model of behavioral data vs. thinking of it as a model of a cognitive system.

joint model is to allow tuning of all of the submodels' parameter weights simultaneously while respecting the constraints imposed on them by the hyperparameters. Since joint models are Bayesian hierarchical models, they impose probabilistic structure across the entire parameter space. The two hyperparameters (the mean vector Φ and the variance-covariance matrix Σ) along with the priors over them determine the multivariate normal distribution from which the first-order parameters are drawn.

The main empirical results center on the comparison of the cognitive model with the bi- and trivariate joint models that incorporate neural data sources. These models were fitted independently and then compared with respect to the distribution of their estimates of the true drift rate parameter values. For the two bivariate models (EEG and fMRI), the variance of the cognitive model parameter estimates was significantly reduced, and the smallest overall variance in parameter estimates was achieved in the trivariate model. This indicates that adding new data sources can improve estimates of model parameters as well as generalization performance on new data. Of the three joint models, the trivariate model achieved the most accurate predictions when used to generate behavioral data for participants whose data was withheld during training. So the improved parameter estimates also come with an advantage in predictive performance.

These conclusions comport with an earlier study by Turner et al. (2015) in which they constructed a similar joint model using only behavioral and fMRI data, and in which the cognitive submodel was the DDM rather than the LBA. For the fitted bivariate joint model, the posterior predictive distributions for the DDM drift rate and starting point parameters had smaller variance than for the cognitive model uncoupled from any neural data. The bivariate model's behavioral predictions also had less variance and it performed better in cross-validation tests that compared its predictions with that of the regular DDM. Both of these improvements are traceable to the model's ability to constrain cognitive model parameters by neural source data. In this case the neural data reduction methods were hypothesis-free: rather than focusing on a prechosen set of ROIs, the whole brain fMRI data was transformed using a series of spatial principal and independent component analyses that settled on 34 independent components as having the most robust and significant activations over all of the trials. This shows that joint model fitting can produce improved parameter estimates even without specific information about relevant regional neural activity.

3.2. Data analysis

A second use of cognitive models is not to refine our understanding of hypothetical cognitive constructs but to analyze neural data. This approach is associated with the research tradition of model-based fMRI (Gläscher and O’Doherty 2010; Love 2020). The aim of this work is broadly exploratory: to find plausible neural correlates for known cognitive constructs defined within the model; or, put in slightly stronger terms, to make tentative assignments of cognitive function to brain regions or networks. This problem of functional assignment is one that drives much of cognitive neuroscience. The method is to take a cognitive model initially fitted to a set of behavioral data and then search for correlations between the model’s latent variables and the neural data. To the degree that the model’s variables are psychologically interpretable, this method may shed light on the functions of particular regions or on the kinds of processes that neural measurements reliably signal (see Turner et al. 2019, 60).

As an example, consider Mulder et al.’s (2012) study of neural signals of decision bias. Participants performed a two-alternative random-dot motion detection task inside an MRI scanner. Before each trial they were given cues (in the form of arrows) that either indicated the likely direction of dot motion, or indicated a higher value (operationalized as monetary payoff) for a direction. The primary psychological question was whether these biasing cues would be effective and if so where in the decision process they have their effect. The DDM was fitted to the response time and accuracy data for each participant in the two bias conditions and the starting point and drift rate parameters of the trained model were compared. The main finding was that for manipulations of both kinds of biasing cues there was a significant increase in starting point, but no significant change in drift rate. This suggests that biasing cues change the state of initial decision making but do not affect how rapidly evidence is processed afterwards.

The second phase of the study investigated where in the brain starting point bias might be encoded and processed. This involved a two-stage analysis of the scans from individual participants’ trials. The first analysis of the neural data used a general linear model (GLM) with 13 regressors corresponding to the different types of cues, stimuli, and feedback the participants received. From the resulting maps, contrast images representing the difference between the bias and neutral conditions were created. The second analysis, which is the crucial one for our purposes, regressed these contrast images in a GLM that used the starting point bias term from

the fitted DDM as a covariate. The overall goal of this processing sequence was to progressively home in on regional BOLD responses that correlated with individual-level differences in bias as revealed by the DDM.

These final analyses returned a number of regions where an increase in starting point correlates with an increase in regional BOLD signal. For the cue that indicates likely direction of motion a set of regions passed this threshold, including right superior frontal gyrus, right middle frontal gyrus, left inferior frontal gyrus, left intraparietal sulcus, medial frontal gyrus, and anterior cingulate gyrus. With some adjustments to the signal threshold, similar regions appeared in the analysis of manipulations targeting the value of a particular direction. The overlap in five frontoparietal regions that showed up in both of these analyses was taken as evidence that they are part of “a common mechanism of bias in choice behavior” (Mulder et al. 2012, 2341), although the ascription of mechanism here is a bit too strong given that the evidence is only correlational.¹⁰ More conservatively, we can conclude that the studies plus the accompanying analysis reveal a cluster of regions that covary with a hypothesized process in early decision-making. Data analysis of this sort is exploratory, not confirmatory. These regions are ones that can serve as targets for interventions and manipulations that may bolster the initial functional claims.

4. Models as interfield evidential integrators

While these sketches cannot convey the full range of methods employed in MBCN, they illustrate how model-building practices can integrate behavioral and neuroscientific data. Integration begins within a context of inquiry that requires deciding how data from these sources can be brought together to answer a particular set of questions. Interfield investigations advance by finding the right model-building toolkit to construct linking relations between data in various

¹⁰ This is also noted by Povich (2019, 11), who claims that MBCN “provides evidence about *what* [neural region] implements the cognitive model components” but not *how* a region implements a function; that is, it can describe structure-function correspondences but is silent on the mechanistic implementation of those relationships.

fields. And, following the pragmatic-representational conception of data, these relations are in turn what make the data evidentially relevant to the questions in play.¹¹

In our first case, the questions at issue concern parameter weights, which are encodings of the relative causal strength of the cognitive constructs at work in simple decision tasks. Joint models incorporating multiple subcomponents are the pivotal mediating structures that allow data to be brought to bear on these questions across fields. Within these models, the evidential relevance of neural data to psychological questions is formally expressed through the hyperparameters of the joint model, and this relevance is partially confirmed by the fact that the fitted joint model outperforms the purely behavioral model in terms of estimated parameter accuracy.

The second case illustrates the same epistemic phenomenon realized by a different route. Here, behavioral data are used to fit a cognitive model that in turn enables us to query a neural data set to explore hypotheses about neural function. There is not a single joint model to which the neural and behavioral submodels both contribute, hence assumptions about evidential relationships cannot be localized to a particular model parameter. Instead we have a *modeling pipeline*—a sequence of interlocking models that feed into each other—whose structure implicitly encodes linking assumptions. Relations between stages of the pipeline and how data is passed along them can encode these assumptions in many ways. For example, assumptions about the correspondences between cognitive constructs and brain activation are implicit in the fact that parameters from a behaviorally-fitted DDM are used as covariates in a sequence of GLMs that progressively sift the neural data for promising leads. Passing these values on from one model to another (possibly with additional intermediary processing) implies linkages among them.

¹¹ Throughout this paper I employ the language of “creating” evidence. This is most consistent with subjectivist views according to which data acquire the status of evidence because we have reasons to take them as confirmationally relevant to some hypothesis. Objectivist views of evidence, by contrast, would interpret this as a situation in which certain pre-existing evidential relations obtain which we only now have reasons to acknowledge. For objectivists, it is the relations, not our acknowledging them, that makes this data evidence. While I have adopted subjectivist language, the account given here is ultimately meant to be neutral between these conceptions.

In all of these cases we have a common factor, namely a cognitive model such as the DDM which can be used to link behavioral and neural data and thereby to establish evidential relations. Call this common factor an *embedded model*. Embedded models are situated within a larger structure which can take either the form of a single overarching model (Section 3.1) or a set of separate models connected via a processing pipeline (Section 3.2). Call the total structure that the embedded model is part of the *embedding model*. A model becomes embedded either by being a submodel of the embedding model, or by being part of the modeling pipeline constituted by the embedding model. Being embedded is therefore a form of mereological containment relation. The epistemic work of building an embedding model lies in choosing the appropriate data sets, a suitable model structure, and well-founded linking relations, as well as experimental, analytic, and computational procedures. Embedding models are tools for uncovering particular sets of facts, but the kind of embedding model built to answer questions about parameter fine-tuning differs from one designed to perform data analysis or assist with the task of selecting among competing models.

There is no recipe or algorithm for determining how to embed a model in a way that makes it adequate for a given purpose, although there are precedents and heuristics that provide guidance. Each new embedding invokes a grab bag of shared tricks, tactics, and tools that can be adjusted and assembled to suit the particular case. The resulting embedding model is a highly specific contraption whose structure is attuned to the materials that went into its construction. The embedded model—the DDM itself—is an invariant and computationally indispensable ingredient in all of these activities, but it is crucial to note that the inferential work of data unification is distributed across the whole apparatus of the embedding structure.

To see this, note that in the hierarchical model (Figure 1), the influences of the two other data sources (EEG and fMRI) are essential to arriving at better distributions of parameter estimates than could be achieved with the DDM taken on its own. It is only the total embedding structure consisting of both data sources, the embedded DDM, and the hyperparameters, that drives these inferences. In our representative modeling pipeline (Figure 2), the DDM feeds parameters to a general linear model that is the product of its own series of transformations, and it is this terminal model that delivers information about what neural data might be of

psychological interest. The embedded DDM only serves as a key source of input values. In both cases it is not the embedded model taken on its own, but only the total embedding model that makes it possible to answer questions within one field by using data from another. Showing how data can be used for the purposes of inquiry is precisely what it is to show the relevance of that data, and thus to establish its evidential value.

These cases also illustrate that embedding models have a hybrid representational structure. They combine an embedded cognitive model with several associated data models originating in different experimental systems (EEG, fMRI, reaction times, and error rates). The component data models represent statistically processed measurements of biophysical and behavioral quantities. Cognitive models, on the other hand, represent causally significant elements of cognitive systems. The fact that cognitive models represent what they do allows them to drive many of the inferences made using embedding models. The whole model (including its explicit and implicit linking propositions) represents the complex of relations among these represented quantities and elements. In the second case study, for example, it is assumed that the starting point bias term represents a particular cognitive quantity. The discovery that this term covaries with measured activity in an fMRI data model underwrites exploratory hypotheses about the functions of the brain regions represented by that activity.

These structures are, therefore, neither purely models of data, nor models of causal structure. Because of their hybrid representational nature, embedding models can serve as inferential engines for integrating data sources produced within diverse experimental systems. This point is particularly worth noting since much of the debate over the nature of models in cognitive neuroscience has centered on how to interpret their mapping to real-world systems. A prominent suggestion due to Kaplan and Craver (2011) is that explanatory models must satisfy the models-to-mechanisms mapping (3M) constraint, which requires that the variables and dependencies within the model correspond to, respectively, components and causal relations in the real-world target system. The 3M constraint is designed to distinguish models that are explanatory from those that are phenomenal. The latter capture the shape of potential explanatory targets, but do not themselves explain them. Versions of the constraint have been explored with respect to dynamical, topological, and network models, among others.

These debates, while significant, still keep the focus on explanation, which is only one epistemic task among many. The present analysis of MBCN shows, however, that the purposes of modeling often demand stitching together field-specific submodels that have different representational functions. Analyzing the kinds of models at work in these cases requires going beyond oppositions such as phenomenal versus explanatory. The role of these models is not to give a formal description of a pattern that serves as target of explanation, nor to capture the causal mechanisms that generate and explain such patterns, even if subcomponents of the overall embedding model may do these things. The function of the embedding model taken as a whole is to construct and explore relations among behavioral and neural data sources and the parameters of cognitive models. Performing this integrative task may, of course, facilitate other epistemic and practical ends (see Section 6), but it is nevertheless a distinct precondition for them.

It is useful to approach the practice of data unification by way of Galison's metaphor of trading zones. Trading zones are, in their original anthropological sense, often literal places where goods and materials are productively exchanged despite having radically different values and meanings for both parties. This is what gives trade its power: it can go on without a neutral currency or even a shared understanding of the act of trading itself (Galison 1997, 803). In a scientific context, trading zones arise where subcultures of a discipline (or distinct fields) share boundaries yet differ in their methods, concepts, objects of study, and values. Galison's most well-documented examples involve the communities of experimentalists, theorists, and instrumentalists in postwar microphysics. The materials exchanged can be instruments, theoretical terms and concepts, computational procedures, and, importantly for our purposes, recordings and other materializations of data. These items are meaningful within each community, but their interpretation is suspended at the point of exchange, enabling them to operate and signify in divergent ways for each participant. Thus it is that "distinct groups, with their different approaches to instruments and their characteristic forms of argumentation, can nonetheless coordinate their approaches around specific practices" (806).

Data unification within MBCN shares many of the properties of epistemic exchanges in Galisonian trading zones. First is their provisional character: data unification is concerned with finding ways to make particular data sets evidentially relevant to specific research questions arising within contexts of inquiry. Once a way has been found to establish these local relevance

relations, the fields can again go their separate ways. Data unification establishes a host of relatively short-duration connections between or among fields. These exchanges enable participants in each field to carry out their own investigations and further their own goals. Provisional interactions do not assume that their participants have common or convergent interests. A psychologist and neuroscientist coordinating their activities may each hope to advance their own field-specific inquiries without concern for whether the others are similarly benefited, or whether they contribute to any common enterprise. Finally, fields may recurrently coordinate with each other on particular questions that arise within each one separately, participating in any number of exchanges without these resulting in any kind of stable overarching group organization. In this way, coordination between fields can be permanently ad hoc, in the sense that there may be no alternative to the practice of building models anew for each specific coordination task.

The asymmetry of data unification means that the benefits of new evidence do not always accrue identically to all participating fields. Neural evidence can refine the accuracy and predictive power of cognitive models, and this in turn allows us to answer questions typically posed in applied and theoretical psychology, but this does not necessarily loop back to inform neuroscientific theory. Similarly, the ability to analyze neural data using cognitive model parameters benefits neuroscientific theorizing and guides future experimentation, but it has no direct and immediate relevance to psychological modeling. Episodes of interfield coordination are often guided by such field-centric concerns. In this sense, unification is not a matter of fields merging or interpenetrating, but of finding a stable region within which necessary epistemic exchanges can be made as needed. When the tools for making such exchanges are more or less ready at hand, fields have achieved a degree of data unification.¹²

¹² In this light, it is unsurprising that textbooks like Turner et al. (2019) are filled with chapters describing and illustrating these techniques. I also note that their title refers specifically to joint modeling of *data*, a point which highlights the way that they themselves view the overarching goal of their project.

5. The practical character of data unification

This overview of how embedding models are constructed highlights the fact that data unification proceeds in parallel with methodological and explanatory unification while nevertheless being importantly distinct from them. MBCN, and data unification more broadly, are best understood as instances of what Grantham (2004) calls *practical* unification. In his taxonomy, practical connections among fields include heuristic dependence, confirmational dependence, and methodological integration. Learning how to trade data with potential evidential value is a precondition for establishing confirmational relations as well as for heuristic purposes such as exploratory analysis. We might view data exchange as a type of interfield connection that is logically prior to others on the list, or perhaps as a type that cross-cuts them. However conceived, it belongs squarely to the practical domain. As Grantham notes concerning analogous cases from other fields, “the question of how (or whether) to use stratigraphic data in phylogeny reconstruction is *not* a problem about how the explanations or theories of paleontology and neontology are related. Rather, it is a problem of resolving potential conflicts between the distinctive data generated by two fields” (148).

The account developed here is a complement to related work on the practical dimensions of data integration. Consider Leonelli’s threefold distinction among interlevel integration, cross-species integration, and translational integration in plant biology (2016, 141–75). She characterizes these forms of integration as “involving the assembling and interrelation of data documenting different levels of organization within the same species... the comparison and co-construction of data on different species... [and] the use of data from a wide variety of sources to devise new forms of intervention on organisms” (143). Interlevel integration is especially pertinent here insofar as it focuses on understanding the same system, e.g., the mouse-ear cress *Arabidopsis thaliana*, from many different “levels” at once, ranging from molecular interactions and cellular metabolic pathways through macroscale phenomena such as the development of root systems and flower morphology. Given that MBCN, too, is concerned with combining data drawn from many different perspectives, it is worth considering her account in greater detail.

Leonelli focuses on the often-overlooked epistemic role played by curatorial practices in crafting databases and repositories, and the embodied and propositional forms of knowledge that underlie the creation of these resources. She offers three examples that illustrate how this form of

database construction works. These can be classified in terms of whether they constitute data production, data interpretation, or data mobilization.

The first two aspects of interlevel integration center on procedures for gathering, sorting, labeling, and disseminating data—that is, they are forms of mobilization. They concern the social, institutional, and epistemic structures that support the construction and maintenance of lasting resources such as databases. First, curation involves the development of metadata codes that mark epistemically significant features of each piece of data logged. Metadata can include facts about what experimental procedures were involved in the data’s production, where and by whom it was produced, and what computational procedures were used to analyze it. From these codes, researchers can determine the fit and reliability of a dataset with respect to their own projects. Second, curation involves creating keywords that enable data to be searched and sorted by researchers from a wide range of backgrounds. This requires consultation with members of many fields to understand how they conceptualize the objects and processes under investigation. The data revolution in science is characterized in large part by the proliferation of these curatorial projects. There are, for example, several databases that collect and make publicly available a vast trove of neural imaging data: the Cognitive Atlas, OpenNeuro, BrainMap, NeuroSynth, and Dev-Atlas, among others.

Sullivan (2017) discusses several of these initiatives, but concludes that they are not (or were not at the time of writing) adequate for the purpose of providing well-grounded taxonomies of psychological constructs and neural functions. The reasons for this deficiency lie in how the datasets are curated and coded. As she observes, “[t]he labels designating psychological functions that current databases contain do not reflect intra and interdisciplinary consensus as to how to generally define or how to produce, detect and measure the phenomena designated by those labels” (134). If this assessment is correct, psychology and neuroscience lag behind the biosciences in developing the institutional, economic, and social infrastructure, as well as the disciplinary incentives, that are necessary for data-centric science to make progress.

However, the account developed here does not focus on the creation and curation of these data repositories. Constructing embedding models is most akin to Leonelli’s final example of data integration. She notes that databases often come with their own software platforms that allow users not just to search and retrieve data, but to merge it: “to combine and visualize

genomic, transcriptomic, and metabolic data as a single body of information” (2016, 147). These are fundamentally interpretive activities. The studies described in Section 3 presuppose that data produced in varying circumstances have already been brought together, by whatever routes. The problem is how to make them evidentially relevant.

Unlike databases and repositories, the particular datasets used in these studies are not (usually) treated as persisting epistemic artifacts. Nor, more significantly, do the embedding models that unify them persist beyond the circumstances of their use. The models constructed within MBCN are inherently ad hoc, transitory things. They are built for a specific purpose to answer a set of locally posed interfield questions and they serve as instruments towards this larger set of goals, rather than being valuable epistemic products in their own right. Joint models, for example, are assembled for data sets with specific characteristics (e.g., how many neural data sources are included, how those sources are processed or reduced, etc.) and need to be redesigned and retrained when these change. Often a substantially different model is the result. The use of an embedding model to construct one set of relations between fields does not necessarily transfer to other contexts, and the model itself is typically set aside once its particular analytic work is done.

Despite the fact that its materials are fleeting, interfield data integration constitutes a stable practice. What persists is a kind of know-how: the ability of participants from several fields to make each other’s data relevant to their own inquiries. This ability draws on a standing body of techniques and tools for building models that reliably bring specific data types to bear on specific research questions. Data integration, then, can go on even in the absence of databases and other archives. It depends on practitioners having the epistemic dispositions to reliably connect information from one field with that from another, as needed. Reliable connectability is underpinned by the existence of data modeling techniques and the knowledge of how to deploy them. Evidential relevance relations themselves can persist so long as the know-how for bringing data to bear in new circumstances does. Establishing a trading zone between two fields involves developing precisely this sort of mutual intelligibility grounded in practices of modeling and interpretation.

6. The autonomy of data unification

The evidence-making epistemic function of models should be distinguished from others that have been floated in discussions of interfield integration and unification. Embedded modeling is, in particular, not a method of explanatory unification. As noted above, these models are not generally explanatory ones. They neither represent targets of explanation nor possible explanations of them. Embedded modeling can be directed at many purposes. These include explanation itself, of course, but also prediction, data exploration, parametric fine-tuning, model selection, and discovery of possible interventions, none of which reduce to explanation. All of this is consistent with their primary role as inferential engines for data unification.

This picture of unification stands in contrast with a view advanced by Nathan (2017). He argues that “of the various interconnections postulated by interfield theories, a single one—explanatory relevance—lies at the core of unifications” (176) and therefore the role of “all other [interfield] relations” is “grounded in, motivated by, and ultimately reducible to their contribution to the explanatory relevance of fields” (177). On this view, methodological and data unification would be subservient to explanatory unification. This proposal requires us to ignore the plurality of autonomous goals involved in scientific research, however. We should particularly distinguish between the proximal and distal goals of building interfield models. The immediate research questions that motivate these constructions are, I’ve argued, not invariably explanatory ones. That is compatible with the fact that establishing that certain kinds of data can be used as evidence may suggest further experiments, constrain the space of live hypotheses, and ultimately lead to (dis)confirmation of explanatory claims. In this very general and long-term sense, the results of these studies may turn out to have explanatory relevance.

Nathan claims that most or all other kinds of integration “can be understood as part of an explanatory endeavor, as long as the notion of explanation is conceived broadly enough” (2017, 177). He does not elaborate on what this broad notion of explanation might look like, however, or how other forms of integration might fall under it. To take one example, sharpening models’ parameter values may be important in designing more targeted psychiatric intervention. Knowing which cognitive parameters are implicated in a disorder can offer important clues for its

treatment. Such translational applications are not themselves explanation-oriented.¹³ My proposal is that these diverse forms of unification have their own point and epistemic texture, and are therefore best thought of as intimately related but ultimately autonomous. The claim of explanatory relevance itself is something that cannot be assumed but only demonstrated through further investigation. Moreover, focusing on potential distal applications pulls our attention away from the other more, immediate concerns that shape evidential exchanges. The examples canvassed here bring out the distinct epistemic structure of data unification and its differences from explanatory and methodological forms of unification.

This perspective on MBCN can usefully be compared with one advanced by Povich. He argues that “[b]y providing mutual constraint from different scientific fields on the mechanisms that underlie and are responsible for cognitive processes, MBCN represents multifield mechanistic integration” (2019, 9). He elaborates that “MBCN realizes (or at least approximates) the mechanistic ideal of explanation by helping to identify potential realizers for the components of cognitive models that are themselves mechanistic explanations of some cognitive capacity” (9). While MBCN does not itself specify mechanisms, it shrinks the possibility space of what structures might realize them. This idea of mutual constraint comes out in the repeated emphasis by MBCN practitioners on using many different data types as evidence for hypotheses in both cognitive psychology and neuroscience. Povich expounds on these constraints by noting that MBCN can provide “suggestive evidence about *what* in the brain realizes” a particular feature of cognition such as decision threshold or recognition strength (10). The study in Section 3.2 exemplified this pattern. Such correlational evidence may serve as a constraint on theorizing about how that capacity is realized mechanistically.

Since mechanisms are paradigmatic explanatory entities, this sounds like an argument that MBCN is a form of explanatory unification after all. However, if we bear in mind the distinction between proximal and distal goals of inquiry, Povich’s claims are compatible with the ones made here. The distal goal of some researchers in these interfield projects might be to

¹³ On this point, see Leonelli’s (2016, 152–58) discussion of translational integration, where she notes that biological research aimed at improving social conditions can and does proceed without shedding light on interlevel and cross-species integration; that is, translation and other forms of understanding can be pursued separately.

describe the mechanisms that implement various elements of cognitive models. Part of achieving this will require establishing mutual constraints between these models and their neural realization base. But establishing these constraints rests in turn on the ability to bring distinct forms of data to bear on the same space of hypotheses, and this is a primary function of model-building in MBCN. The process of data unification is autonomous in the following sense: first, that it has a separate purpose from explanatory unification; second, that it has a distinct epistemic structure from explanatory unification; and third, that it can be carried out not just prior to but also without any long-term commitment to explanatory unification.

7. Conclusions

Given the shift in recent decades to ever more data-centric scientific practices, it is widely recognized that we need to understand how forms of data generated in distinct fields can be brought to bear on one another. But the problem of unification long precedes this shift. Overcoming it has been the constant struggle of the cognitive sciences since their inception—as, for example, the tortuous history of psycholinguistics will testify.

The account of data unification presented here highlights the fact that these modes of epistemic trade often turn on model-building practices. Data journeys and other modes of data unification have been characterized elsewhere in great detail. The focus in this discussion has instead been on the ways that models play key roles in creating evidence. Coordinating scientific activity across fields requires knowing how to craft embedding models that can successfully do the inferential work needed to establish the relevance of data to new kinds of questions, and thereby turn it into evidence. A virtue of MBCN is that it provides exemplary cases of how attempts at making data evidentially relevant might succeed. At the same time, these integration projects are not invariably successful (O'Malley 2013). Further research will have to clarify the potential failure modes of these practices, the scenarios in which data unification fails, and whether certain kinds of data simply cannot be integrated with each other at all.

References

- Alexandrowicz, Rainer W. 2020. "The Diffusion Model Visualizer: An Interactive Tool to Understand the Diffusion Model Parameters." *Psychological Research* 84 (4): 1157–65. <https://doi.org/10.1007/s00426-018-1112-6>.
- Alexandrowicz, Rainer W., and Bartosz Gula. 2020. "Comparing Eight Parameter Estimation Methods for the Ratcliff Diffusion Model Using Free Software." *Frontiers in Psychology* 11 (September):484737. <https://doi.org/10.3389/fpsyg.2020.484737>.
- Arnold, Nina R., Arndt Bröder, and Ute J. Bayen. 2015. "Empirical Validation of the Diffusion Model for Recognition Memory and a Comparison of Parameter-Estimation Methods." *Psychological Research* 79 (5): 882–98. <https://doi.org/10.1007/s00426-014-0608-y>.
- Bechtel, William. Forthcoming. "The Epistemology of Evidence in Cognitive Neuroscience." In *Philosophy and the Life Sciences: A Reader*, edited by Robert A. Skipper, Colin Allen, Rachel A. Ankeny, Carl F. Craver, Lindley Darden, Gregory Mikkelsen, and Robert C. Richardson. Cambridge: MIT Press. <http://mechanism.ucsd.edu/epist.evidence.bechtel.july2004.pdf>.
- Bokulich, Alisa, and Wendy Parker. 2021. "Data Models, Representation and Adequacy-for-Purpose." *European Journal for Philosophy of Science* 11 (1): 31. <https://doi.org/10.1007/s13194-020-00345-2>.
- Darden, Lindley. 2005. "Relations among Fields: Mendelian, Cytological and Molecular Mechanisms." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 36 (2): 349–71. <https://doi.org/10.1016/j.shpsc.2005.03.007>.
- Darden, Lindley, and Nancy Maull. 1977. "Interfield Theories." *Philosophy of Science* 44 (1): 43–64. <https://doi.org/10.1086/288723>.
- Forstmann, Birte U., Roger Ratcliff, and Eric-Jan Wagenmakers. 2016. "Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions." *Annual Review of Psychology* 67 (1): 641–66. <https://doi.org/10.1146/annurev-psych-122414-033645>.
- Galison, Peter. 1997. *Image and Logic: A Material Culture of Microphysics*. Chicago: University Of Chicago Press.

- Gläscher, Jan P., and John P. O’Doherty. 2010. “Model-Based Approaches to Neuroimaging: Combining Reinforcement Learning Theory with fMRI Data.” *Wiley Interdisciplinary Reviews: Cognitive Science* 1 (4): 501–10. <https://doi.org/10.1002/wcs.57>.
- Grantham, Todd A. 2004. “Conceptualizing the (Dis)Unity of Science.” *Philosophy of Science* 71 (2): 133–55. <https://doi.org/10.1086/383008>.
- Grasman, Raoul P.P.P., Eric-Jan Wagenmakers, and Han L.J. van der Maas. 2009. “On the Mean and Variance of Response Times under the Diffusion Model with an Application to Parameter Estimation.” *Journal of Mathematical Psychology* 53 (2): 55–68. <https://doi.org/10.1016/j.jmp.2009.01.006>.
- Hollander, Gilles de, Birte U. Forstmann, and Scott D. Brown. 2016. “Different Ways of Linking Behavioral and Neural Data via Computational Cognitive Models.” *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* 1 (2): 101–9. <https://doi.org/10.1016/j.bpsc.2015.11.004>.
- Irvine, Elizabeth. 2016. “Model-Based Theorizing in Cognitive Neuroscience.” *The British Journal for the Philosophy of Science* 67 (1): 143–68. <https://doi.org/10.1093/bjps/axu034>.
- Kaplan, David Michael, and Carl F. Craver. 2011. “The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective.” *Philosophy of Science* 78 (Oct.): 601–27. <https://doi.org/10.1086/661755>.
- Leonelli, Sabina. 2013. “Integrating Data to Acquire New Knowledge: Three Modes of Integration in Plant Science.” *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 44 (4): 503–14. <https://doi.org/10.1016/j.shpsc.2013.03.020>.
- Leonelli, Sabina. 2016. *Data-Centric Biology: A Philosophical Study*. Chicago ; London: The University of Chicago Press.
- Leonelli, Sabina. 2019. “What Distinguishes Data from Models?” *European Journal for Philosophy of Science* 9 (2): 22. <https://doi.org/10.1007/s13194-018-0246-0>.
- Lerche, Veronika, and Andreas Voss. 2017. “Retest Reliability of the Parameters of the Ratcliff Diffusion Model.” *Psychological Research* 81 (3): 629–52. <https://doi.org/10.1007/s00426-016-0770-5>.

- Love, Bradley C. 2020. "Model-Based fMRI Analysis of Memory." *Current Opinion in Behavioral Sciences* 32 (April):88–93. <https://doi.org/10.1016/j.cobeha.2020.02.012>.
- Mulder, Martijn J., Eric-Jan Wagenmakers, Roger Ratcliff, Wouter Boekel, and Birte U. Forstmann. 2012. "Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff." *Journal of Neuroscience* 32 (7): 2335–43. <https://doi.org/10.1523/JNEUROSCI.4156-11.2012>.
- Nathan, Marco J. 2017. "Unificatory Explanation." *British Journal for the Philosophy of Science* 68 (1): 163–86. <https://doi.org/10.1093/bjps/axv022>.
- O'Malley, Maureen A. 2013. "When Integration Fails: Prokaryote Phylogeny and the Tree of Life." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 44 (4): 551–62. <https://doi.org/10.1016/j.shpsc.2012.10.003>.
- O'Malley, Maureen A., and Orkun S. Soyer. 2012. "The Roles of Integration in Molecular Systems Biology." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 43 (1): 58–68. <https://doi.org/10.1016/j.shpsc.2011.10.006>.
- Povich, Mark. 2019. "Model-Based Cognitive Neuroscience: Multifield Mechanistic Integration in Practice." *Theory & Psychology* 29 (5): 640–56. <https://doi.org/10.1177/0959354319863880>.
- Ratcliff, Roger. 1978. "A Theory of Memory Retrieval." *Psychological Review* 85 (2): 59–108. <https://doi.org/10.1037/0033-295X.85.2.59>.
- Ratcliff, Roger, and Russ Childers. 2015. "Individual Differences and Fitting Methods for the Two-Choice Diffusion Model of Decision Making." *Decision* 2 (4): 237–79. <https://doi.org/10.1037/dec0000030>.
- Ratcliff, Roger, Philip L. Smith, Scott D. Brown, and Gail McKoon. 2016. "Diffusion Decision Model: Current Issues and History." *Trends in Cognitive Sciences* 20 (4): 260–81. <https://doi.org/10.1016/j.tics.2016.01.007>.
- Shiffrin, Richard M., Michael D. Lee, Woojae Kim, and Eric-Jan Wagenmakers. 2008. "A Survey of Model Evaluation Approaches with a Tutorial on Hierarchical Bayesian Methods." *Cognitive Science* 32 (8): 1248–84. <https://doi.org/10.1080/03640210802414826>.

- Sullivan, Jacqueline Anne. 2017. "Coordinated Pluralism as a Means to Facilitate Integrative Taxonomies of Cognition." *Philosophical Explorations* 20 (2): 129–45. <https://doi.org/10.1080/13869795.2017.1312497>.
- Teller, Davida Y. 1984. "Linking Propositions." *Vision Research* 24 (10): 1233–46. [https://doi.org/10.1016/0042-6989\(84\)90178-0](https://doi.org/10.1016/0042-6989(84)90178-0).
- Turner, Brandon M., Birte U. Forstmann, Bradley C. Love, Thomas J. Palmeri, and Leendert Van Maanen. 2017. "Approaches to Analysis in Model-Based Cognitive Neuroscience." *Journal of Mathematical Psychology* 76 (February):65–79. <https://doi.org/10.1016/j.jmp.2016.01.001>.
- Turner, Brandon M., Birte U. Forstmann, and Mark Steyvers. 2019. *Joint Models of Neural and Behavioral Data*. Computational Approaches to Cognition and Perception. Cham, Switzerland: Springer Nature Switzerland AG. <https://doi.org/10.1007/978-3-030-03688-1>.
- Turner, Brandon M., Birte U. Forstmann, Eric-Jan Wagenmakers, Scott D. Brown, Per B. Sederberg, and Mark Steyvers. 2013. "A Bayesian Framework for Simultaneously Modeling Neural and Behavioral Data." *NeuroImage* 72:193–206. <https://doi.org/10.1016/j.neuroimage.2013.01.048>.
- Turner, Brandon M., Leendert van Maanen, and Birte U. Forstmann. 2015. "Informing Cognitive Abstractions through Neuroimaging: The Neural Drift Diffusion Model." *Psychological Review* 122 (2): 312–36. <https://doi.org/10.1037/a0038894>.
- Turner, Brandon M., Christian A. Rodriguez, Tony M. Norcia, Samuel M. McClure, and Mark Steyvers. 2016. "Why More Is Better: Simultaneous Modeling of EEG, fMRI, and Behavioral Data." *NeuroImage* 128 (March):96–115. <https://doi.org/10.1016/j.neuroimage.2015.12.030>.
- Woodward, James. 1989. "Data and Phenomena." *Synthese* 79 (3): 393–472. <https://doi.org/10.1007/BF00869282>.
- Woodward, James. 2000. "Data, Phenomena, and Reliability." *Philosophy of Science* 67:S163–79. <https://doi.org/10.1086/392817>.
- Woodward, James. 2010. "Data, Phenomena, Signal, and Noise." *Philosophy of Science* 77 (5): 792–803. <https://doi.org/10.1086/656554>.




Figure 1. Hierarchical model illustrating one form of model embedding. Hyperparameters are pictured at center, and arrows represent relations by which subparameters are set by the hyperparameters.

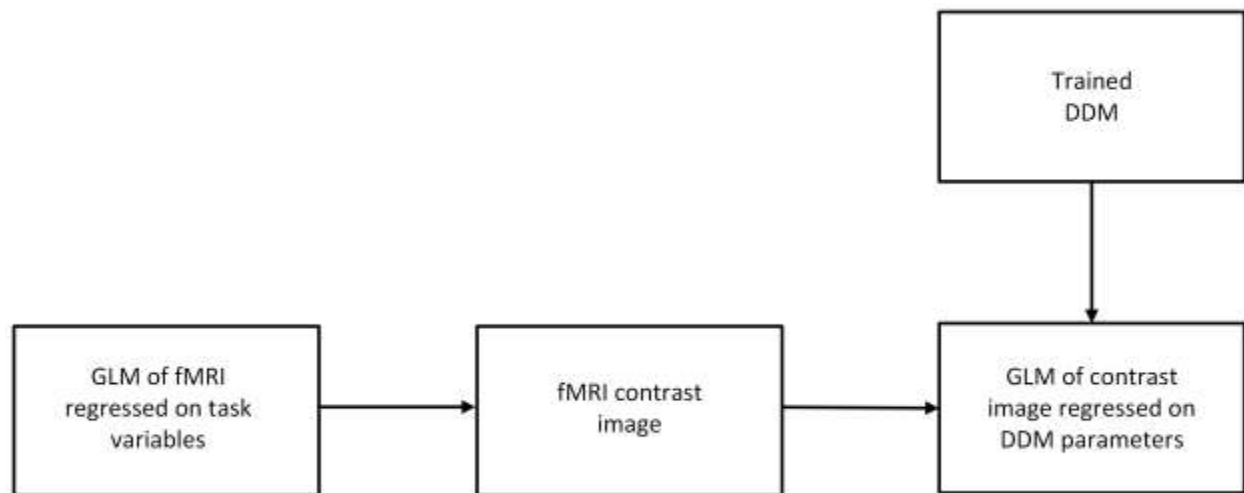


Figure 2. Modeling pipeline illustrating another form of model embedding. Arrows represent either model transformations or ways of passing parameters from one model to another.