Deep learning for robust meta-analytic estimation

Evan C. Carter

U.S. Army Research Laboratory, Aberdeen, MD, USA

Author Note

Correspondence concerning this article should be addressed to Evan Carter, Email: evan.c.carter@gmail.com or evan.c.carter3.civ@mail.mil. A set of scripts to run the re-trained version of DeepMA on users' own data is available at <https://osf.io/a72rf/>. This is a working paper that is actively being submitted for peer-review. This version was last modified on March 20, 2019. Future versions may include substantial changes.

Abstract

Meta-analysis represents the promise of cumulative science—that each successive study brings us greater understanding of a given phenomenon. As such, meta-analyses are highly influential and gaining in popularity. However, there are well-known threats to the validity of meta-analytic results, such as processes like publication bias and questionable research practices which can cause researchers to massively overestimate the evidence in support of a claim. There are many statistical methods to correct for such bias, but no single method has been found to be robust in all realistic conditions. Here, I describe a method that merges statistical simulation and deep learning to achieve an unprecedented level of robust meta-analytic estimation in the face of numerous forms of bias and other historically problematic conditions. Furthermore, the resulting estimator, called DeepMA, has the unique property that it can easily evolve: As new conditions for which robustness is needed are identified, DeepMA can be re-trained to maintain high performance. Given the weaknesses that have been identified for meta-analysis, the current consensus is that it should serve as simply another data point, rather than residing at the top of the hierarchy of evidence. The novel approach I describe, however, holds the potential to eliminate these weaknesses, possibly solidifying meta-analysis as the platinum standard in scientific debate.

Keywords: meta-analysis, publication bias, questionable research practices, machine learning, deep learning, neural networks.

Deep learning for robust meta-analytic estimation

Introduction

A study in an empirical science such as Medicine, Psychology, Economics, or Physics, only provides a limited glimpse of the phenomena it is designed to assess. The purpose of meta-analysis—which is the statistical aggregation of results from a set of similar primary studies—is to magnify and clarify this glimpse. Meta-analyses, therefore, can strongly guide both policy and the direction of future research; however, despite being highly influential and growing in popularity (de Vrieze, 2018), meta-analysis has some well documented—and sometimes severe—limitations.

Probably the most discussed limitation is that meta-analytic findings can be severely biased if the results of a primary study determine whether or not it is included in the meta-analytic sample. For example, pre-conceived notions about which direction the evidence in a field points can influence how a meta-analyst designs inclusion criteria, thereby guaranteeing that the meta-analysis confirms an existing opinion. Indeed, if a meta-analyst is willing to capitalize on so-called “researcher degrees of freedom,” nearly any conclusion can be reached (de Vrieze, 2018). This is bias implemented at the level of the meta-analysis and, fortunately, can be addressed through incentivizing rigorous transparency (e.g., pre-registration, disclosure of conflicts of interest) and tactics such as adversarial collaboration (Kahneman, 2003). In other words, these are problems that are ultimately solvable through cultural shifts in how meta-analysts perform their work.

Processes that introduce bias at the level of the study, however, are a much harder problem to address. For example, primary studies with statistically non-significant or counter-theory results may never be made public (i.e., lost in the “file drawer”; Rosenthal, 1979), thereby skewing the meta-analytic sample. As a result, so-called publication bias can create the appearance of strong meta-analytic evidence when in fact there is none. Critically, since most meta-analyses rely on an existing literature, positive cultural changes implemented now cannot retroactively reduce bias in studies that were conducted before the meta-analysis was ever conceived. Instead, meta-analysts rely on statistical methods to correct for study-level bias.

Unfortunately, the general consensus is that none of the existing state-of-the-art methods perform reliably well across all realistic conditions (e.g., Carter, Schönbrodt, Gervais, and Hilgard, 2017), and because meta-analysts can never know the exact conditions that generated the data they work with, it may be impossible to draw conclusions when estimators yield diverging results. A truly robust meta-analytic estimator is therefore badly needed. Here, I describe a novel approach that combines meta-analysis, statistical simulation, and deep learning to produce exactly such a method. In the following, I present both a new estimator that represents a dramatic improvement upon the state-of-the-art, as well as an iterative process that can improve this estimator as our understanding of the limits of meta-analysis grows.

Meta-analysis, simulation, and machine learning.

Meta-analysis is typically formulated as follows. The meta-analytic sample is made up of $i = 1, \dots, k$ independent effect size estimates d_i . It is assumed that $d_i = \delta + e_i$, where δ is the true underlying effect estimated by each d_i . e_i is the associated sampling error, and $e_i \sim \mathcal{N}(0, v_i)$, where v_i is the sampling variance and is assumed known. The fixed-effect meta-analytic model takes these assumptions and provides $\hat{\delta}$, an estimate of the true underlying effect measured by each of the k primary studies (Hedges & Vevea, 1998).

Because most primary studies in a meta-analytic sample are not identical, the true underlying effect is often modeled as the mean of a distribution of true underlying effects. In that case, the i^{th} study measures a study-specific true underlying effect $T_i = \delta + u_i$, where $u_i \sim \mathcal{N}(0, \tau^2)$. Thus, τ^2 represents the variance in true underlying effects that might reflect the influence of study-level characteristics such as idiosyncratic methods or the demographics of the sample. The random-effects meta-analytic model takes these assumptions and provides $\hat{\delta}$ and $\hat{\tau}$ as estimates of the average true underlying effect and the degree of between-study heterogeneity, respectively.

The estimates produced by these models can be severely biased by processes like publication bias or the undisclosed use of researcher degrees of freedom (also called

questionable research practices, or QRPs; John, Loewenstein, and Prelec, 2012) during production of the primary studies (Rosenthal, 1979; Moreno et al., 2009; Carter et al., 2017). The extent to which either of these processes influence a literature is, by definition, unknown. But enough evidence has emerged (e.g., John et al., 2012; Dwan et al., 2008; Franco, Malhotra, and Simonovits, 2014; Fanelli, Costas, and Ioannidis, 2017) to cast doubt on all but the most carefully monitored efforts (e.g., large-scale registered replication reports; Simons, Holcombe, and Spellman, 2014). As a result, many bias-correcting estimators have been proposed, and many simulation studies have been reported to demonstrate the usefulness of such estimators (e.g., Carter et al., 2017; Moreno et al., 2009; Stanley and Doucouliagos, 2014; Simonsohn, Nelson, and Simmons, 2014; McShane, Böckenholt, and Hansen, 2016; R. C. van Aert, Wicherts, and van Assen, 2016).

For a bias-correcting estimator to be truly useful, it must be robust—that is, it must perform well across simulated conditions that plausibly represent the (unknowable) processes that generate real-world data. Among other things, simulation studies usually assess whether a bias-corrected estimate $\hat{\delta}$ remains close to the known value of the true underlying effect δ , regardless of factors such as the magnitude of δ , the size of the meta-analytic sample k , the size of between-study heterogeneity τ^2 , or the degree to which bias is introduced by processes like publication bias. Other factors, such as the distributional assumptions underlying between-study heterogeneity (normal or uniform, for example) or the effect of QRPs, are less frequently investigated, but must also be considered. As mentioned, to date, simulation studies provide little reason for researchers to feel confident that any state-of-the-art method is robust (Carter et al., 2017). From the perspective of those looking for clear-cut answers from meta-analysis, this is a very disappointing result; however, there are two active areas of research in the field of machine learning that may provide the key to moving the field of meta-analysis forward.

First, the fact that simulation studies show that most meta-analytic methods have different strengths and weaknesses suggests that such methods make up a set of diverse

learners—that is, estimators that individually only perform well in certain conditions. Critically, research in machine learning has found that when diverse learners are combined into an ensemble, the ensemble can provide far more robust estimation (Brown, Wyatt, Harris, & Yao, 2005). Second, because meta-analysis—and the difficulties it faces can be encapsulated by formally-defined simulations, one can draw inspiration from research areas where simulation and machine learning are used in concert to powerful effect. For example, standard algorithms for complex but tractable problems, such as the calculation of the volume displaced by a moving object, can be limited in their usefulness because of their computational demand. However, what is known about such problems can be implemented as a simulation, which can then be used to produce massive amounts of data for training deep neural networks. Those trained networks can then be used in place of the standard algorithms at a dramatically reduced computational cost and with minimal loss in accuracy Chiang, Faust, and Tapia, 2018.

The work I present here draws directly on these two themes in machine learning research to develop a truly robust meta-analytic estimator. The first step is to identify the conditions one wishes the resulting estimator to be robust to, and then program them into a simulation. Here, I focused on varying forms and degrees of publication bias, QRPs, between-study heterogeneity, primary and meta-analytic sample sizes, and values of the true underlying effect δ . The second step is to train machine learning algorithms using data produced by the simulation to return the known parameters of interest—in this case, the values of δ and τ that were used in the simulation to produce the data.

Although any supervised machine learning algorithm can be used for this approach, here I focused on deep neural networks that take 118 input features, including variables such as the estimates from ten existing meta-analytic methods, and return values for $\hat{\delta}$ and $\hat{\tau}$. As an abbreviation for “deep learning for robust meta-analytic estimation,” I refer to the resulting estimator as DeepMA.

Methods

A brief description of deep learning

The kind of deep learning used to develop DeepMA is a sub-topic within the larger area of supervised machine learning. In this framework, a machine learning algorithm is “trained” on many observed pairs of input and output data, and what is learned are the “rules” that relate inputs to outputs. Learning is done by minimizing the difference between the predicted output and the observed output relative to some loss function (e.g., mean squared error). The predicted output can be thought of as a transformation of the input data, so the process is one of learning useful “representations” of the input data.

The “deep” in deep learning refers to the fact that this approach is focused on learning successive layers of representations of the input data, where the layers are models called neural networks. Each layer consists of a number of units (sometimes called neurons). In the first layer, each unit can correspond to an element of input data, what is often called a dependent variable or predictor in other statistical applications. Each unit in each layer after the first does some form of computation that combines the values of the units it is connected to in the previous layer with a weight parameter to produce an output. For example, for fully connected layers, the output of units in the first layer $\mathbf{X}$ are combined with the weights of the units in the second layer $\mathbf{W}$ such that the signals associated with the second layer are $\mathbf{Z} = \mathbf{XW}$. The signals are then transformed—through, for example, a sigmoid function—to become the output of the second layer, and thus, the input to the third layer. This process repeats up to the final layer, which provides a predicted output. The predicted output is then compared to the known output using a loss function, and the resulting loss is used to adjust the weights in each layer in a way that is likely to decrease it.

This adjustment is done by some form of mini-batch stochastic gradient descent, referred to as an optimizer. Essentially, a randomly selected batch of training data is run through the network to determine the loss, then the gradient of the loss is calculated with respect to the weights of the network. The weights are then updated in

the opposite direction of the gradient to minimize the loss for the batch. This is repeated for some number of iterations known as epochs. At the start of training, weights are randomly initialized, and the predicted output will likely be far from the known output (i.e., a high loss). As training progresses, weight values that minimize the loss function will be found.

Training. Eventually, the network can be considered “trained” in that the predicted outputs are close to the observed outputs. However, because machine learning algorithms are notoriously good at over-fitting training data, training must be done in a way that ensures out-of-sample generalizability (Chollet, 2017). Typically, this is achieved by conducting training with a subset of the total data—a so-called “training set”—and then evaluating performance using data not in the training set—what is usually called a “test set.” Good performance on the test set can be taken as evidence that a network has learned a representation that is not unique to the training set.

In addition to the weight parameters of a network, tuning of hyperparameters also needs to be done. Hyperparameters are any feature of the network other than the weights: For example, the architecture of the network (i.e., number of layers and the number of units in each layer, as well as the connections between the units in each layer), the optimizer used for learning, the activation function for each layer, and the size of the mini-batches and epochs. Tuning these parameters is the act of finding the values for them that result in the best performance on a particular problem—in other words, it is another form of learning. Therefore, just like the tuning of the network’s weight parameters during training, tuning hyperparameters needs to be done carefully to avoid over-fitting the network to the training data. The usual solution to this is to tune hyperparameters using performance on the “validation set” (a subset of the data that is distinct from both the training and test sets) as feedback.

After developing a group of candidate models with different hyperparameter configurations that show reasonable performance on the validation data, one can select a final model based on which of the candidate models in the group shows the best performance on the test set. In principle, this will be the network that has the most

generalizable performance to new data.

Transfer learning. The point of training a network as described above is to learn representations that are generalizable between the training, validation, and test sets, with the assumption that this will make generalizability to new data more likely. However, it is also important to consider generalizability to new data that have been produced by different processes than those that generated the data in the training, validation, and test sets. This form of generalizability is sometimes called transfer learning—when what is learned in one context can be transferred to another (Pan, Yang, et al., 2010).

Transfer learning is highly relevant to DeepMA because of two assumptions that it must make. First, that good estimation of known parameters used to produce simulated data (e.g., the average true effect, δ) is highly positively correlated with good estimation of unobservable parameters underlying real world data¹. In other words, that what is learned on simulated data will transfer to real world data.

The second assumption required by DeepMA is that a version of DeepMA trained on data generated from a finite set of simulation conditions will show transfer learning to new data generated from different simulation conditions. This assumption is required because there are an infinite number of possible simulation conditions, and it is therefore impossible to expose DeepMA to each one during training. Instead, the best that can be done is to generate training data from sufficiently representative conditions. For example, publication bias can, and has, been implemented in simulations in various ways (e.g., Copas and Shi, 2000; Moreno et al., 2009; Carter et al., 2017). If one implementation produced data patterns that were representative of all of the others, then DeepMA would, in principle, learn enough from being trained only on data from that one implementation to perform well on data from all of the others. Developing a useful version of DeepMA, therefore, involves testing for transfer learning using data generated from new simulation conditions.

Fortunately, whenever DeepMA fails to show transfer learning on data from a new

---

¹ This same assumption is implied when simulation studies are used to assess other meta-analytic methods, so it is not unique to DeepMA.

condition, it can be re-trained using those new data. The approach I describe here, therefore, does not end with the versions of DeepMA I present below. Instead, as criticisms or questions regarding the conditions used to train the current versions arise, DeepMA can—with relative ease—evolve. This advantage is unique to DeepMA.

Datasets

In the current work, I make use of two datasets to demonstrate both transfer learning and the re-training process that allows DeepMA to evolve. The first dataset, which I refer to as the “Initial Dataset” is based on data from Carter et al., 2017, which was generated from a simulation framework designed to encompass publication bias and QRPs, as well as data from a slightly modified version of that simulation. The second dataset, which I refer to as the “Transfer Dataset,” consists of data generated from a simulation involving further modifications to the framework used by Carter et al., 2017. All simulations were conducted in R (R Core Team, 2016).

initial Dataset. As mentioned, part of the initial dataset are the data described by Carter et al., 2017. The code for that simulation is available at <https://osf.io/rf3ys>, and a detailed description of it is provided by Carter et al., 2017. I briefly describe it again in the following.

Each simulated meta-analytic sample of k simulated primary studies was associated with an average true underlying effect δ and a true underlying degree of between-study heterogeneity τ . For the i^{th} simulated primary study, n independent samples were drawn for both the control and experimental group— y_c and y_e , respectively—where $y_c \sim \mathcal{N}(0, 1)$, $y_e \sim \mathcal{N}(T_i, 1)$, and $T_i \sim \mathcal{N}(\delta, \tau^2)$. The i^{th} simulated primary study was therefore represented by a tuple including d_i as an estimate of T_i (in the form of a standardized mean difference effect size; Cohen’s d), the variance v_i associated with the effect size d_i , the per-group sample size n , and the p -value and t -statistic associated with a t -test comparing y_c and y_e .

The group sample size n was drawn randomly from an inverse gamma distribution (truncated below $n = 5$ and above $n = 1905$), the parameters ($a = 1.153$ and $b = 0.046$)

of which were determined by finding the closest fit to real group sample sizes drawn from studies reported in several major psychology journals from 1995 to 2006 (Marszalek, Barber, Kohlhart, & Cooper, 2011).

When data were affected by publication bias, the result from a simulated primary study entered into the final meta-analytic sample with probability w , which was mapped to the simulated study's one-tailed p -value by one of two selection functions representing either medium publication bias or high publication bias. Both the medium and high publication bias functions set $w = 1$ when $p_{one-tailed} < 0.025$ (i.e., when it was statistically significant and in the theory-consistent direction). Both functions also set w to exponentially decrease from $p_{one-tailed} = 0.025$ to $p_{one-tailed} = 0.05$ and then to stay at a constant level until $p_{one-tailed} = 0.50$. From $p_{one-tailed} > 0.50$ to $p_{one-tailed} = 0.995$, $w = 0$, and then from $p_{one-tailed} = 0.995$ to $p_{one-tailed} = 0.9995$, w increased exponentially until reaching a constant level at $p_{one-tailed} > 0.9995$. The rate of exponential change was faster and the constant level lower for high publication bias as compared to medium publication bias.

The patterns described above formalize the idea that theory-consistent studies stand some chance of being published even if they are not statistically significant (i.e., $w > 0$), and that theory-inconsistent studies can be published if they are "highly" statistically significant. These chances are less in the presence of high publication bias as compared to medium publication bias.

In addition to publication bias, data could also be affected by four independent QRPs: 1) Optional removal of outliers (i.e., data points that were more than two standard deviations from the mean), 2) optional selection between two dependent variables (which were correlated at $r = 0.20$), 3) optional use of moderators, and 4) optional stopping. The exact implementation of these is described in detail by Carter et al., 2017.

Simulated researchers could take three strategies: 1) no use of QRPs (the "pure" strategy); 2) the use of optional dependent variables and some degree of optional stopping (adding three more observations to both y_e and y_c and re-checking results up

to three times) during data collection (the “moderate” strategy); and 3) the use of optional outliers, optional dependent variables, optional moderators, and a higher degree of optional stopping during data collection (adding three more observations to both y_e and y_c and re-checking results up to five times; the “aggressive” strategy).

Each simulated meta-analytic sample occurred under one of three QRP environments, which determined the mixture of QRP strategies that would be used for a simulated study: 1) none (all studies produced under the pure strategy); 2) medium (30% of studies produced with the pure strategy, 50% with the moderate strategy, and 20% with the aggressive strategy); and 3) high (10% of studies produced with the pure strategy, 40% with the moderate strategy, and 50% with the aggressive strategy).

QRPs were applied before publication bias—that is, a simulated researcher using QRPs attempted to get their results to the point where $p_{one-tailed} < 0.025$, but if they could not, the result could still be kept out of the meta-analytic sample by publication bias. Simulated studies were generated until k observations were entered into the meta-analytic sample.

The simulation data from Carter et al., [2017](#) includes 432,000 simulated meta-analyses: 1000 simulated meta-analyses for each of the 432 conditions made up of every unique combination of $k = \{10, 30, 60, 100\}$, $\delta = \{0, 0.2, 0.5, 0.8\}$, $\tau = \{0, 0.2, 0.4\}$, publication bias = $\{none, medium, high\}$, and QRP environments = $\{none, medium, high\}$. These data were supplemented with an additional 420,000 simulated meta-analyses from the same framework as described above, but where the parameters were determined independently for each simulated meta-analysis with random draws from probability distributions. For this supplemental data, the parameters were distributed as follows: $k \sim \mathcal{U}(10, 200)$, $\delta \sim \exp(\lambda = 3)$, and $\tau \sim \exp(\lambda = 3)$, and each form of publication bias and QRP environment had an equal probability of being chosen. For half of these supplemental simulated meta-analyses (i.e., 210,000), all signed values (i.e., δ , d , t) were multiplied by -1 . Therefore, these supplemental data represented negative true underlying effects and versions of both publication bias and QRPs that preferred results where $d < 0$, in contrast to the

original focus on only positive values by Carter et al., 2017.

The total 852,000 simulated meta-analyses that made up the Initial Dataset were randomly divided into a training set (678,615 simulated meta-analyses), a validation set (88,183 simulated meta-analyses), and a test set (85,202 simulated meta-analyses). The full initial dataset is available at <https://osf.io/a72rf/>.

Transfer Dataset. The simulation used to generate the Transfer Dataset involved further modifications to the framework of Carter et al., 2017. Instead of simulating the true underlying effect for a simulated meta-analysis such that $T_i \sim \mathcal{N}(\delta, \tau^2)$, T_i was simulated as following either a uniform or gamma distribution: $T_i \sim \mathcal{U}(a, b)$ or $T_i \sim \text{gamma}(a, b)$. For the case of uniformly distributed heterogeneity, the parameters were $a = \frac{2\delta - \sqrt{12}\tau^2}{2}$ and $b = \sqrt{12}\tau^2 - \frac{2\delta - \sqrt{12}\tau^2}{2}$. For the gamma distributed heterogeneity case, the parameters were $a = \frac{\delta}{\tau^2/\delta}$ and $b = \frac{\tau^2}{\delta}$. Additionally, instead of drawing per-group sample sizes n from an inverse gamma distribution, n was equally likely to follow either $n \sim \mathcal{U}(3, 15)$ or $n \sim \mathcal{U}(150, 250)$. In this way, more extreme sample sizes were simulated in the transfer dataset than in the initial dataset.

In addition to the above, the implementations of publication bias and QRPs were also modified. For high publication bias, the function relating the probability w to the one-tailed p -value was simplified so that only a study with statistically significant and theory-consistent results would enter the meta-analytic sample (i.e., $w = 1$ when $p_{\text{one-tailed}} \leq 0.025$, $w = 0$ otherwise). For QRPs, the high QRP environment was modified so that 100% of the studies were created using the aggressive strategy. Moreover, the aggressive QRP strategy was made more intense by increasing the amount of data simulated researchers could collect during optional stopping: Two more observations for both y_c and y_e were collected up to fifteen times. Additionally, the correlation between dependent variables was set to $r = 0$.

The transfer simulation also involved different parameter values: $k = \{5, 8, 125, 250\}$, $\delta = \{0.1, 0.5\}$, $\tau = \{0.01, 0.8\}$, publication bias = $\{\text{none}, \text{high}\}$, and QRP environments = $\{\text{none}, \text{high}\}$. For each of the 128 conditions created by each unique combination of these parameter values, 2000 simulated meta-analyses were

created for a total of 256,000 simulated meta-analyses. The 256,000 simulated meta-analyses in the transfer dataset were randomly divided into a training set (207,360), a validation set (23,040), and a test set (25,600). The full transfer dataset is available at <https://osf.io/a72rf/>.

Meta-analytic methods

DeepMA takes as input the results from ten existing meta-analytic estimators: the random-effects meta-analysis model (Hedges & Vevea, 1998); the trim-and-fill (Duval & Tweedie, 2000); the precision effect test (PET; Stanley and Doucouliagos, 2014); the precision effect estimate with standard error (PEESE; Stanley and Doucouliagos, 2014); the conditional combination of PET and PEESE (PET-PEESE; Stanley and Doucouliagos, 2014); the weighted average of the adequately powered (WAAP-WLS; J. Ioannidis, Stanley, and Doucouliagos, 2017; Stanley, Doucouliagos, and Ioannidis, 2017); the three and four parameter selection models (McShane et al., 2016; Iyengar & Greenhouse, 1988); p -uniform (R. C. van Aert et al., 2016); and p -curve (Simonsohn et al., 2014). These methods are described in detail in the work cited above and more briefly in Carter et al., 2017. The exact implementation of these methods (using R; R Core Team, 2016) can be found on the OSF page associated with Carter et al., 2017: <https://osf.io/rf3ys>. In the following, I briefly describe the conceptual basis for each method.

The random-effects meta-analysis model and the trim-and-fill method.

The random-effects model assumes that each individual primary study measures some true underlying effect T_i , and that the true underlying effects can be considered to have been drawn from a normal distribution (i.e., $T_i \sim \mathcal{N}(\delta, \tau^2)$). The model provides an estimate of both the mean δ and standard deviation τ of this distribution. The model is essentially an average of the study-level effect size estimates d weighted by the inverse of the effect size variances (i.e., $1/v$), and thus, offers no correction for bias.

The trim-and-fill is an extension of the random-effects model and is probably the most commonly implemented bias-correcting method. The trim-and-fill is based on the

funnel plot, which is a scatter plot where effect sizes d are plotted against some proxy for sample size (e.g., standard error $\sqrt{v}$). In the presence of bias—or any process that creates a correlation between effect size estimates and sample sizes, also known as small-study effects; Sterne, Gavaghan, and Egger, 2000—the funnel plot loses its characteristic symmetrical funnel shape. The trim-and-fill works by first testing for so-called funnel-plot asymmetry, and when it exists, iteratively “trimming” observations until the funnel is symmetrical. It then “fills” back in the trimmed observations along with imputed versions of each on the opposite side of the center of the funnel. This new meta-analytic sample that contains imputed observations is then re-analyzed with a naive method such as the random-effects model.

Here, the trim-and-fill and the random-effects model were implemented using the default options of the R package `metafor` (Viechtbauer, 2010).

Meta-regression. WAAP-WLS, PET, PEESE, and PET-PEESE are all referred to as meta-regression methods because they are all based on a weighted-least squares (WLS) regression model with inverse-variance weights (just like the random-effects meta-analysis model).

The rationale for WAAP-WLS follows from the premise that estimates from primary studies with adequate power should be preferred over any other estimates, in part because the influence of bias is reduced by statistical power (J. Ioannidis et al., 2017; Stanley et al., 2017). To conduct WAAP-WLS, first, an intercept-only WLS regression model is run on the meta-analytic sample to get an estimate of the true underlying effect. Then, a subsample of only those studies that show at least 80% power to detect that estimate is defined. Next, another intercept-only WLS model is run on only this new subsample of observations (i.e., the adequately powered studies). The estimate of the true underlying effect from this second model is the output of WAAP-WLS.

Both PET and PEESE, like the trim-and-fill, aim to correct for small-study effects. PET is a WLS regression model where effect size is a function of the standard error of effect size. The intercept of this model can be considered an estimate of the

true underlying effect that has been corrected for any influence of small-study effects. PEESE is essentially the same, but instead of regressing effect size on the standard error of effect size, effect size is regressed on standard error squared (i.e., variance). The interpretation of the intercept in PEESE is the same as in PET, but PEESE differs in that effect size is a curvilinear function of the standard error as opposed to a linear one.

PET-PEESE is a conditional combination of the estimates from PET and PEESE. Simulation studies (Stanley & Doucouliagos, 2014) have shown that PET is a better estimator when the true underlying effect is zero, but that PEESE is better otherwise. Thus, the output of PET-PEESE is the estimate from PET when that estimate is statistically non-significant (relative to a one-tailed p -value in the direction of the theorized effect) and the estimate from PEESE otherwise.

All of these methods were implemented in R exactly as they were implemented by Carter et al., 2017.

Selection models and methods based on p -value distributions. Selection models are similar to the random-effects meta-analysis model in that they allow estimation of parameters representing the average true underlying effect (δ) and between-study heterogeneity (τ), but they also include parameters that represent the probability of certain studies entering the observed meta-analytic sample. For example, in addition to the parameters δ and τ , the three-parameter selection model implemented here also includes the parameter p_1 , which represents the probability of a study with a one-tailed p -value < 0.025 (i.e., in the expected direction and statistically significant) entering the sample (this is the model developed by Iyengar and Greenhouse, 1988 and discussed recently by McShane et al., 2016).

The four-parameter selection model implemented here, in addition to δ , τ and p_1 , also includes p_2 , which represents the probability of studies with one-tailed p -value > 0.50 entering the sample (in other words, studies in the counter-theory direction).

For both the three- and four-parameter selection models, maximum-likelihood estimation is used to provide parameter estimates. Following Carter et al., 2017, both of these models are implemented using the `weightr` package (Coburn & Vevea, 2017).

As recommended, when there are less than four observed p -values in each bin created by the probability parameters, no estimates are produced.

Two recent methods— p -curve (Simonsohn et al., 2014) and p -uniform (R. C. van Aert et al., 2016)—can be understood as single parameter selection models that only estimate δ . These models assume $\tau = 0$ and that only statistically significant studies enter in to the sample (so there is no need to estimate probability parameters). These two methods differ from one another in terms of the estimation method used, and so in general, should be expected to provide the same results.

Carter et al., 2017 implemented p -curve using R scripts based on those supplied by Uri Simonsohn and following the authors’ recommendations: 1) Only statistically significant and directionally-consistent studies were supplied to p -curve, 2) the p -curve estimate was set to 0 when no eligible studies were available or when p -curve returned an estimate in the unexpected direction.

Carter et al., 2017 implemented p -uniform using the R package `puniform` (van Aert, 2017) and followed the authors’ recommendation that the estimate of the true underlying effect was set to zero when the average p -value in the sample was less than 0.025.

Both p -curve and p -uniform provide additional information beyond the estimate of δ . For p -curve, there are also statistical tests associated with three alternative hypotheses: 1) that the distribution of statistically significant p -values shows right-skew (i.e., that there is evidential value in the data); 2) that the evidential value in the data is inadequate; and 3) that the distribution of statistically significant p -values shows left-skew, indicating the presence of bias.

For p -uniform, there is also a test for the presence of publication bias based on the degree to which the distribution of statistically significant p -values (conditional on the assumption that the true underlying effect equals the effect size estimate from a fixed-effect meta-analysis model) deviates from a uniform distribution. It is worth noting that the authors of p -curve do not interpret their method’s estimate of δ to be an estimate of the true underlying effect. Instead, the estimate from p -curve should be

taken as the estimate of the average of the effects supplied to it if publication bias was reduced to zero (Simonsohn et al., 2014). Thus, the results presented here for p -curve should be interpreted with this in mind, since p -curve was not designed to be assessed in this way.

Inputs to DeepMA. Each simulated meta-analysis was represented in terms of an input vector with 120 elements, or features: the meta-analytic sample size k (1 feature); the estimate of δ from each of the ten methods described above (10 features); the standard error of that estimate for all methods except p -uniform and p -curve (8 features); the test statistic, p -value, and upper and lower 95% confidence limits associated with the estimate of δ for all methods except p -curve (36 features); the I^2 value and Q -statistic from the random-effects meta-analysis model (standard metrics of between-study heterogeneity; 2 features); the estimate of τ and the standard error of that estimate from the random-effects meta-analysis model, trim-and-fill, and the three- and four-parameter selection models (8 features); the test statistic, p -value, and upper and lower 95% confidence limits associated with the estimate of τ from the three- and four-parameter selection models (8 features); the estimate, test statistic, p -value, and upper and lower 95% confidence limits of the weight parameters in the two selection models (note that the three-parameter selection model has one of these and the four-parameter selection model has two; 18 features); the estimate, test statistic, p -value, and upper and lower 95% confidence limits for the non-intercept parameter in PET, PEESE, and PET-PEESE (18 features); the number of adequately powered studies identified by WAAP-WLS (1 feature); the number of studies imputed into the funnel plot by the trim-and-fill (1 feature); the test statistic and p -value associated with p -curve’s three additional hypothesis tests (6 features); the test statistic and p -values from the test for publication bias provided by p -uniform (2 features); and the number of statistically significant studies in the meta-analytic sample (1 feature). Any instance for any of the above features that took a value of NA or infinity was set to zero.

Each feature for every simulated meta-analysis was then centered and scaled relative to those in the initial training set. Every pairwise correlation for the 120

features in the initial training set was also calculated, and any feature showing a correlation of $|r| > 0.999$ with another feature was removed. The features removed from the initial training set were also removed from the other five sets, although the correlations between features in these sets were never calculated. Two features were removed through this process: the standard error of the estimate of δ by WAAP-WLS and the test statistic for p -curve’s test of evidential value. The result was 118 final features. The above pre-processing was done using the R package `caret` (Kuhn, 2008).

DeepMA development

Framework for training and hyperparameter tuning. All training was done in python using Keras (<https://keras.io/>; Chollet et al., 2018) as a frontend to TensorFlow (www.tensorflow.org; Abadi et al., 2016), the widely used Google library for deep learning. Hyperparameter optimization was automated using Hyperas (github.com/maxpumperla/hyperas), a convenience wrapper for the Hyperopt library (github.com/hyperopt/hyperopt; J. Bergstra, Yamins, and Cox, 2013). This process involves tree-structured Parzen Estimator approach (J. S. Bergstra, Bardenet, Bengio, & Kégl, 2011) and can be more efficient than manual or random hyperparameter tuning.

I first designed six candidate network architectures (Figure 1). Each architecture started with an input layer of 118 units (one unit per 118 features described above), contained some number of fully connected “hidden” layers, and terminated in two output units, one to provide $\hat{\delta}$ and one to provide $\hat{\tau}$. The activation function for the $\hat{\delta}$ unit was linear and the activation function for $\hat{\tau}$ unit was softplus (so estimates of τ were constrained to be ≥ 0). Loss associated with $\hat{\delta}$ was assigned twice the weight as loss associated with $\hat{\tau}$. This weighting scheme reflects my *a priori* belief that the primary goal of meta-analysis is to estimate the true underlying effect δ as well as the *post hoc* observation that during initial exploratory analyses, τ was generally easier to estimate.

During training, I defined loss as mean absolute error (MAE). For the j^{th} simulated meta-analytic dataset, each estimator provides an estimate $\hat{\delta}_j$ of the true

underlying effect δ_j . MAE for that estimator a set of M simulated meta-analytic datasets is therefore $\frac{\sum_{j=1}^M |\hat{\delta}_j - \delta_j|}{M}$. MAE is the average mis-estimation of the true underlying effect, so meta-analytic methods with lower MAE are to be preferred. Although root mean squared error (RMSE; see below) is a more typical metric with which to evaluate the performance of meta-analytic methods in simulation studies, I chose to train DeepMA using MAE because MAE is linear in that estimation errors are weighted equally, regardless of their distance from the true value. The implication of this property is that no error is necessarily better or worse than any other, which in my view, is true of meta-analysis. For example, if $\delta = 0$, $\hat{\delta} = 0.02$ is not necessarily twice as bad as $\hat{\delta} = 0.01$. Instead, I argue that error should be judged in the context of the research topic (e.g., in terms of the practical significance of a given degree of misestimation; Carter et al., 2017), and this assessment is more easily conducted with MAE.

For each of the six candidate architectures, Hyperas was used to tune the following hyperparameters to minimize MAE on the initial validation set: Units per layer = {5, 10, 25, 50, 100, 200, 250, 300, 400}; activation type per layer = {the rectified Linear Unit, the hyperbolic tangent activation function}; optimizer = {RMSProp, Adam, stochastic gradient descent, Adagrad}; mini batch size = {50, 200, 800, 1200, 2000, 2500}; and number of epochs = {50, 100, 200, 300, 400}. Note that each layer could differ in terms of number of units and activation type, whereas optimizer type, mini batch size, and number of epochs was defined for the entire network.

Once trained, these six networks were all then applied to the initial test set *and* the transfer test set, and the network with the lowest MAE for δ on the initial test set was taken as the initial version of DeepMA. No further training or hyperparameter tuning was conducted after the performance of these six networks on the initial test set was observed.

To re-train DeepMA, the same process as described above was followed, but the training, validation, and test sets from the initial and transfer datasets were combined

into new combination training, validation, and test sets. The network that performed best on the combination test set was taken as the re-trained version of DeepMA. Again, no further training or hyperparameter tuning was conducted after the performance of these six models on the combined test set was observed.

Assessing estimation performance

Estimation performance for each method was examined for both the initial test set and the transfer test set, as well as for either all of the two test sets or four subsets of the test sets: 1) no bias and low heterogeneity ($\tau < 0.10$); 2) no bias and high heterogeneity ($\tau > 0.40$); 3) high bias and low heterogeneity; and 4) high bias and high heterogeneity. To be considered robust, a method must show good estimation performance across the four subsets. To demonstrate transfer learning, a method must perform well on both the initial and transfer test sets.

It is difficult to develop an absolute definition of good estimation performance in the context of meta-analytic simulations. However, for the sake of clearly communicating results, I defined the threshold of good estimation performance as the worst estimation performance of the random-effects meta-analysis model in subsets of the initial test set with no bias (i.e., essentially the conditions it was designed for). If a method shows estimation performance at or below this threshold, it can be considered to show good estimation performance.

In addition to considering an absolute definition of good estimation performance, it is also important to consider the performance of an estimator relative to all other estimators. To make this comparison easier, I used a ranking scheme in which ranks represented performance relative to the method with the best performance for a given subset. The method that showed the best performance was assigned rank 1, then the next best performing method was compared to it using the Wilcoxin rank sum test on the distributions of error from the two methods.² If this test was not statistically significant, then the next best performing method was also assigned rank 1. If the test

---

² The Wilcoxin rank sum test was applied to the distributions of absolute values of error $|\hat{\delta} - \delta|$ for MAE and bias, and the distributions of squared errors $(\hat{\delta} - \delta)^2$ for RMSE.

was statistically significant, it was assigned rank 2. This process continued until all methods had been assigned a rank.

Additional analyses of estimation by DeepMA

In addition to assessing the initial and re-trained versions of DeepMA in terms of MAE relative to δ , I also examined estimation performance from several other perspectives.

Estimating τ . I examined estimation performance relative to τ in terms of MAE and following the same procedure as for performance relative to δ . Note that only four of the ten existing meta-analytic methods provide an estimate of τ .

With small k . Many meta-analytic samples have relatively few studies in them. This may be due to the fact that the topic area of interest is new or relatively inactive, or it may be because a large literature on a topic has been divided into smaller subsamples. Regardless, it is important that any meta-analytic method be robust to small meta-analytic sample sizes. The performance for DeepMA and the other estimators was assessed in meta-analytic samples where $k \leq 10$. Note that in the initial test set, 13% of simulated meta-analyses had $k \leq 10$. In the transfer test set, 50% of simulated meta-analyses had $k \leq 10$. In both test sets, the minimum value of k was 5.

With small δ . The influence of bias and heterogeneity on mis-estimation can be stronger for meta-analyses that measure smaller true underlying effects (e.g., Carter et al., 2017). I therefore assessed the performance of DeepMA and the other estimators on meta-analytic samples where $\delta \leq 0.20$. In the initial test set, 47.6% of simulated meta-analyses reflected a value of $\delta \leq 0.20$. In the transfer test set, 50% of simulated meta-analyses reflected a value of $\delta \leq 0.20$.

Mean error (ME) and root mean squared error (RMSE). In addition to MAE, it is reasonable to assess DeepMA's estimation performance in terms of mean error (ME), also called bias, and root mean squared error (RMSE).

Using the same notation as described above for MAE, ME for the j^{th} estimator on a set of M simulated meta-analytic datasets is $\frac{\sum_{j=1}^M \hat{\delta}_j - \delta_j}{M}$. Notice that ME can be both positive (representing over-estimation of the true underlying effect) or negative

(representing under-estimation), but values closest to zero are to be preferred. RMSE is given as $\sqrt{\frac{\sum_{j=1}^M (\hat{\delta}_j - \delta_j)^2}{M}}$. Notice that, like MAE, RMSE is only positive.

It is important to examine the ME of an estimator because metrics such as MAE and RMSE can be lower by virtue of the estimator making consistent mistakes—imagine a distribution of predictions that is tightly centered on the wrong value. However, it is also important to examine a metric such as MAE or RMSE because ME can be lower by virtue of making symmetric but extreme mistakes—imagine a wide distribution of estimates centered on the correct value.

Results

The most complex candidate architecture—number six in Figure 1—outperformed all other architectures in the initial test set, so the initial version of DeepMA took this form. The best performing architecture when re-training DeepMA was one of the least complex candidates—number two in Figure 1—so the re-trained version took this form. A simpler architecture with fewer parameters is harder to “over-train” to a dataset, so the shift from the most complex architecture to one of the least might reflect the increase in variety of the initial training set as compared to the transfer training set.

The condition information for initial and transfer test sets (e.g., δ , τ , k) and the associated values for $\hat{\delta}$ and $\hat{\tau}$ produced by both the initial and re-trained versions of DeepMA have been made available at <https://osf.io/a72rf/>. These data can be used by readers to more fully explore the results (e.g., by examining performance in different subsets).

A set of scripts to run the re-trained version of DeepMA on users’ own data is available at <https://osf.io/a72rf/>.

MAE for the initial version of DeepMA

Figure 2 shows MAE performance for the initial version of DeepMA and the other ten estimators for both the initial test set and the transfer test set, as well as the four subsets described above. I defined the threshold of good estimation performance as the

worst performance of the random-effects meta-analysis model on the initial test set and in the absence of bias: $\text{MAE} \leq 0.07$.

For the initial test set, DeepMA showed good estimation performance in all cases except the high-heterogeneity/high-bias subset. In this subset, DeepMA exceeded the threshold by only $\text{MAE} = 0.01$. Furthermore, when assessed relative to the other estimators, DeepMA performed exceptionally well: For example, it was always ranked first and even outperformed the random-effects meta-analysis model on the no-bias subsets. These results are strong evidence for DeepMA’s robustness.

DeepMA did not perform as well on the transfer test set. For example, it exceeded the threshold for good estimation performance in the two high-heterogeneity subsets, and did so to a fairly high degree in the high-bias/high-heterogeneity subset. However, in both of these subsets, it was still ranked first. In the two low-heterogeneity subsets of the transfer test set, DeepMA was ranked second; however, it still showed good estimation performance, and was never especially outperformed. Given DeepMA’s first-place performance on the entire transfer test set, it seems quite clear that a degree of transfer learning was achieved.

Unfortunately, DeepMA did not maintain an acceptable degree of robustness to high heterogeneity in the transfer test set. Because one of the major differences between the initial and transfer datasets was how heterogeneity was simulated, it is possible that the representations the initial version of DeepMA learned for heterogeneity were simply too limited to be of use on data in the transfer test set. However, as is clear from Figure 3, the re-trained version of DeepMA has solved this problem. Indeed, the re-trained version of DeepMA is ranked first in every subset and shows good estimation performance in every subset except the high-bias/high-heterogeneity subset in the initial test set—here, it exceeds the threshold by only $\text{MAE} = 0.02$.

These results demonstrate exactly the kind of robustness one would hope for in a meta-analytic method, and given the range of conditions on which DeepMA showed good performance, it seems likely that DeepMA’s performance will generalize to data from other simulation conditions that have not yet been examined.

Estimating τ . Figure 4 shows MAE performance for the initial version of DeepMA and the other four estimators that provide estimates of τ for both the initial test set and the transfer test set, as well as the four subsets described above. Based on the random-effects meta-analysis model, the threshold for good estimation performance in this context is defined by $\text{MAE} \leq 0.06$.

The initial version of DeepMA was always ranked first when estimating τ on the initial test set and showed consistently good estimation performance in all subsets (i.e., demonstrated robust estimation performance). However, it performed substantially worse on the high-heterogeneity subsets of the transfer test set (though not much worse than any other method). These findings indicate that insufficient transfer learning occurred. As with estimation of δ by the initial version of DeepMA, it seems likely that the dramatic shift in how heterogeneity was simulated in the transfer test set as compared to the initial test set caused this loss in performance.

Again, the re-trained version of DeepMA dramatically reduced these weaknesses (Figure 5). The re-trained version of DeepMA was ranked first in all subsets of the transfer test set and goes from showing MAE of 0.40 and 0.34 in the high- and no-bias subsets with high heterogeneity, respectively, to showing MAE of 0.10 and 0.07. This exceeds the threshold for good estimation performance, but only by (at most) $\text{MAE} = 0.04$. Moreover, DeepMA’s performance is substantially better than that of any of the other methods: The next best performing methods in these subsets show MAE of 0.35 and 0.28, which exceed the threshold for good estimation performance by (at most) $\text{MAE} = 0.29$.

Notably, the re-trained version of DeepMA showed worse performance on the initial test set as compared to the initial version of DeepMA: For example, it went from being ranked first in all subsets of the initial test set to being ranked second and third in the low-heterogeneity subsets. This behavior is not especially surprising given the fact that the re-trained version of DeepMA must estimate values for τ that come from very different data generation processes. Of course, the drop in performance shown by the re-trained version of DeepMA on the initial test set was very small—difference in

MAE of approximately 0.01—and it still showed good estimation performance. This slight decrease in estimation performance therefore seems like a worthwhile trade-off for the dramatic improvement in performance on the transfer test set.

MAE when $k \leq 10$ or $\delta \leq 0.20$. Figures 6 and 7 show MAE for both the initial and re-trained versions of DeepMA across both the initial and transfer test sets where $k \leq 10$. Figures 8 and 9 show MAE for both the initial and re-trained versions of DeepMA across both the initial and transfer test sets where $\delta \leq 0.20$.

In general, the same patterns as seen across all values of k and δ in the initial and transfer test sets are also seen at values of $k \leq 10$ or $\delta \leq 0.20$ (Figure 10): Namely, the initial and re-trained versions of DeepMA perform very similarly on the initial test set, but the initial version of DeepMA does not perform as well on the transfer test set as the re-trained version of DeepMA. This difference was most extreme in the face of high heterogeneity ($\tau \geq 0.40$) and high bias. The only time this general pattern was deviated from was at $\delta \leq 0.20$ and $\tau \leq 0.10$ in the presence of high bias. In this case, both the initial and re-trained versions of DeepMA performed slightly better on the transfer test set as compared to the initial test set.

Across all conditions examined, the re-trained version of DeepMA never showed performance worse than MAE = 0.12 (at high bias, $\tau \geq 0.40$, and $k \leq 0.10$ k 10). Additionally, the re-trained version of DeepMA was always ranked first in all subsets of the test data examined, and in some cases outperformed the next best method by a difference in MAE of > 0.20 . Moreover, the re-trained version of DeepMA, in every case, showed good estimation performance (see the dashed lines in Figure 10 which represent the worst estimation performance of the random-effects model on the initial test set and in the absence of bias). These results indicate that the dominance shown by DeepMA is not specifically linked to particular values of k or δ , and that the re-trained version of DeepMA provides good estimation performance across all conditions—that is, DeepMA is robust to small meta-analytic sample sizes and small true underlying effects.

Mean error (ME) and root mean squared error (RMSE). Figures 11 and 12 show performance by both the initial and re-trained versions of DeepMA on both the

initial and transfer test sets in terms of mean error. Figures 13 and 14 show performance by both the initial and re-trained versions of DeepMA on both the initial and transfer test sets in terms of RMSE.

Performance by both the initial and re-trained versions of DeepMA was broadly similar when assessed in terms of ME or RMSE as compared to when assessed in terms of MAE. For example, the initial version of DeepMA showed good estimation performance on all subsets of the initial test set and tended to be ranked first. When DeepMA was outperformed, it was in conditions where all other methods did exceptionally well and did not represent a substantial loss in performance by DeepMA.

There was also evidence for a useful degree of transfer learning in that the initial version of DeepMA was ranked first on the entire transfer test set when assessed in terms of both ME and RMSE and that it maintained good performance in all subsets except the high-heterogeneity/high-bias subset.

Just as when performance was assessed with MAE, the re-trained version of DeepMA reduced or eliminated the weaknesses of the initial version. Across all subsets in both the initial and transfer test sets, the re-trained version of DeepMA was off from first place by (at most) a difference of ME of approximately 0.02 and a difference of RMSE of approximately 0.01. Moreover, when the re-trained version of DeepMA exceeded the threshold for good estimation performance, it did so by (at most) ME of approximately 0.02 and RMSE of 0.05. These findings indicate that assessing DeepMA in terms of ME or RMSE does not change the conclusion that it is a more robust estimator than any of the other methods examined, and that the variety of conditions on which it performed well suggest that it will generalize well to new data.

Discussion

Limitations and Future directions

Hypothesis testing and prediction distributions. Many researchers think in terms of null hypothesis significance testing (i.e., whether the null hypothesis can be rejected) instead of in terms of estimation, and therefore expect meta-analysis to

address this question. However, the current version of DeepMA is not designed to test for the rejection of the null hypothesis.

In part, I made this choice because the literal definition of the null hypothesis (i.e., $\delta = 0$) seems unlikely to exist in reality (Cohen, 1990). Furthermore, if the ultimate goal of meta-analysis is to guide policy and future research, then a dichotomous conclusion regarding whether an effect is zero or not seems less useful than a point estimate that is as accurate as possible. For example, if one wants to conduct a power analysis for the purposes of designing a future study, knowing the approximate size of the effect is necessary, whereas simply knowing that it is non-zero provides no guidance. Furthermore, knowing the approximate size of the effect encourages one to consider questions of practical significance. For example, it may be that $\delta = 0.15$, which, although greater than zero, may not be practically meaningful in terms of some real-world medical treatment, therapy, or experimental design.

However, although DeepMA was intentionally made without a focus on null hypothesis significance testing, it is still possible to test whether the null is rejected because DeepMA provides an estimate of τ that can be used to construct a confidence interval around the estimate of δ . This is done by making the same assumptions as in random-effects meta-analysis, defining some critical value C_α (e.g., 1.96 given $\alpha = 0.05$ and assuming a normal distribution), and then calculating $v \cdot = \frac{1}{\sum_{i=1}^k 1/v_i^*}$, which is the standard error of the estimate of δ . As in random-effects meta-analysis, $v_i^* = \hat{\tau}^2 + v_i$, where v_i is the variance of the effect size estimate for the i^{th} study. However, unlike in standard meta-analysis, instead of calculating $\hat{\tau}^2$ through the usual methods (e.g., Viechtbauer, 2005), it can be calculated from the value of $\hat{\tau}$ provided by DeepMA. The upper and lower confidence intervals for the value of $\hat{\delta}$ provided by DeepMA can then be defined as $\hat{\delta} \pm C_\alpha \times \sqrt{v \cdot}$. If zero falls within the bounds defined by these upper and lower limits, then one can say that the null hypothesis has not been rejected.

In addition to a confidence-interval-based approach to null hypothesis significance testing, an entirely new deep learning model can be constructed to classify a given meta-analytic sample as associated with either $\delta = 0$ or $\delta \neq 0$. This approach follows

the same process as described above for DeepMA, however, the ground truth is defined as a boolean representing $\delta = 0$ and the loss function must be one that is appropriate for binary classification (e.g., binary cross-entropy) instead of MAE. The advantage of this method over the above confidence-interval-based one is that it makes full use of the strengths of deep learning instead of making potentially unrealistic assumptions (e.g., that between-study heterogeneity is normally distributed). Despite this, it still focuses on the likely unrealistic null hypothesis of $\delta = 0$ and it still fails to provide the kind of nuance one gets from a point estimate.

Although I believe a meta-analytic point estimate is more useful than a meta-analytic test of the null hypothesis, the point estimate is still limited in that it does not provide information on uncertainty. Developing methods that quantify uncertainty is an active area of machine learning: For example, in the case of deep learning, researchers have proposed using drop out (Gal & Ghahramani, 2016) or ensembles of networks (Lakshminarayanan, Pritzel, & Blundell, 2017) to create prediction distributions, and methods for fully Bayesian neural nets are becoming easier to implement (Tran et al., 2016). Another method that seems particularly promising builds on the ensembling approach to create an approximate Bayesian posterior distribution (Pearce, Zaki, Brintrup, & Neel, 2018). It would be extremely useful to build future versions of DeepMA that provide distributions of estimates rather than only point estimates, since this would add information regarding uncertainty to the estimation of δ and τ .

Adding new features to DeepMA. As mentioned above, a specific set of 118 features were used in developing both the initial and the re-trained versions of DeepMA. However, there is no reason to limit future versions of DeepMA to these features. One of the easiest ways to do this would be to make use of other meta-analytic estimators. For example, limit meta-analysis (Rücker, Schwarzer, Carpenter, Binder, & Schumacher, 2010) and the Copas selection model (Copas, 1999), the recently improved version of p -uniform (called p -uniform*; R. van Aert and van Assen, 2018), and alternative implementations (Moreno et al., 2009) of many of the meta-regression

methods used by Carter et al., 2017. Recently, researchers have also proposed a Bayesian implementation of a selection model called BALM (Du, Liu, & Wang, 2017) that could provide other novel features to DeepMA.

In addition to other estimators, methods that test for the existence of bias could be introduced. For example, the test for excess significance (J. P. Ioannidis and Trikalinos, 2007; 44) and the z -curve (Brunner & Schimmack, 2018). It would also be possible to add fit statistics for each estimator: For example, one could calculate the Akaike Information Criterion (Akaike, 1974) for the standard methods and meta-regression methods. Ultimately, the more features that can be engineered and tested as part of DeepMA, the greater its potential. This should therefore be an active area of future research.

Meta-analytic moderators. DeepMA is designed to provide point estimates for δ and τ because such estimates are a typical goal of meta-analysis. However, another common reason to do meta-analysis is to explore how study-level variables moderate the true underlying effect of interest—for example, does a treatment work particularly well in one population as compared to another. The only way to do this with the current version of DeepMA is through splitting the meta-analytic sample into sub-samples based on a moderator and then estimating δ for each sub-sample. In principle, these estimates could then be compared to examine the patterns of moderation.

There are two drawbacks to this approach. First, if the moderator is continuous, one would have to discretize it, a process that can likely be done in many ways and so can lead to spurious results. Second, a sub-sample is necessarily smaller than the full sample, and it is not fully clear how well DeepMA will perform with very small samples. At the very least, as described above, DeepMA’s performance was robust to small meta-analytic samples sizes in the initial and transfer test sets ($5 \leq k \leq 10$; Figure 7), so this approach seems viable.

The best way to handle meta-analytic moderators is to develop a new version of DeepMA that is specifically designed to do so. For example, just as the simulation used to produce the initial and training datasets in this work was based loosely on the

random-effects model, future datasets could be produced by simulations based on the mixed-effects model. In that case, instead of simply estimating the average true underlying effect δ , where $T_i = \delta + u_i$, one would estimate the parameters (e.g., β_0) of a linear model that specified how the true underlying effect related to p meta-analytic moderator variables x , such that $T_i = \beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip} + u_i$.

New simulation conditions. There are several simulation conditions that are not addressed by the versions of DeepMA described above, but that are likely good candidates for future work aimed at assessing transfer learning.

Processes such as publication bias, QRPs, and heterogeneity can be modeled in many different ways, and it is difficult to argue that any particular implementation is more realistic than any other. Different versions of these processes, therefore, should be a major focus for future research on the limits (and then improvement) of DeepMA’s robustness. For example, one can model publication bias as focused solely on two-sided p -values and ignore the direction of the estimated effects. Additionally, other researchers have simulated publication bias using the process assumed by the Copas selection model (Rücker et al., 2010) or by simply removing some percentage of the most unfavorable results (Moreno et al., 2009). Likewise, QRPs have been modeled in different ways (e.g., “ambitious p -hacking”; Simonsohn, Simmons, and Nelson, 2015), and, in principle, heterogeneity can take on nearly any distributional form.

Beyond publication bias, QRPs, and heterogeneity, future work should also address aspects of the simulated primary data. For example, the use of different effect size metrics. The methods and results discussed here are all focused on meta-analysis of standardized mean difference effect sizes (i.e., Cohen’s d). Although this is an extremely common metric across many scientific literatures, there are several other effect size metrics that are also popular (e.g., log-odds, correlation coefficients). In general, it is possible to transform any one of these effect size metrics into any other (using, for example, the formulas given by Fritz, Morris, and Richler, 2012), so meta-analysts likely do not need to be concerned that DeepMA will be inappropriate for data that take a different form (so long as they transform it to standardized mean differences before

analyzing it with DeepMA). Moreover, because publication bias and QRPs are generally considered to focus on statistical significance and the direction of the effect, and because heterogeneity is independent of the metric of effect size (it is a characteristic of the distribution of true effect sizes), it seems unlikely that DeepMA’s robust performance in the face of bias and heterogeneity will fail to generalize to data in the form of different effect size metrics.

Additionally, the simulations used to train DeepMA assume that the primary study outcome data— y_c and y_e —are normally distributed, so it could be that DeepMA’s performance degrades as this assumption is violated. However, this seems unlikely given that the t -test, which can be considered a transformation of the standardized mean difference or other effect size metrics, has been shown to be robust to such violations (Boneau, 1960). Regardless, this issue, as well as the issue of different effect size metrics, can be addressed empirically, and should be the focus of future research.

Conclusion

Here, I describe an approach that combines meta-analysis, simulation, and deep learning to produce a meta-analytic estimator—DeepMA—that provides significantly more robust estimation than any of the other state-of-the-art methods examined and that can evolve robustness to new conditions as they are conceived.

The core limitation of DeepMA is the limitation of simulations in general: The real world is certainly more complex than any simulation and there is no way to guarantee that a given simulation accurately reflects real processes. The concern, then, is that if real data are generated by a process that is sufficiently different from the one represented by the simulation that produced the data DeepMA was trained on, DeepMA may not perform well on those real data. This concern, however, should be relieved by the fact that DeepMA’s performance generalizes to data produced by simulations that differ from those that produced the data it was trained on (Figure 2), as well as the fact that, so long as a process can be simulated, DeepMA has the capacity to learn to perform well on data produced by that process Figure 3.

Due to the influence of bias at the meta-analytic and study levels described above, a growing number of researchers have argued that meta-analysis should not necessarily reside at the top of the hierarchy of evidence (Stegenga, [2011](#); de Vrieze, [2018](#)). However, the possibility of using deep learning for robust meta-analytic estimation—in conjunction with increased transparency and collaboration at the level of the meta-analysis—revives the potential of meta-analysis to function as the platinum standard in scientific debate.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., . . . Isard, M., et al. (2016). Tensorflow: A system for large-scale machine learning. In *Osdi* (Vol. 16, pp. 265–283).
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6), 716–723.
- Bergstra, J. S., Bardenet, R., Bengio, Y., & Kégl, B. (2011). Algorithms for hyper-parameter optimization. In *Advances in neural information processing systems* (pp. 2546–2554).
- Bergstra, J., Yamins, D., & Cox, D. D. (2013). Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms. In *Proceedings of the 12th python in science conference* (pp. 13–20). Citeseer.
- Boneau, C. A. (1960). The effects of violations of assumptions underlying the t test. *Psychological bulletin*, 57(1), 49.
- Brown, G., Wyatt, J., Harris, R., & Yao, X. (2005). Diversity creation methods: A survey and categorisation. *Information Fusion*, 6(1), 5–20.
- Brunner, J. & Schimmack, U. (2018). Estimating population mean power under conditions of heterogeneity and selection for significance. *Submitted for publication*: <http://www.utstat.utoronto.ca/brunner/papers/Zcurve6.7.pdf>.
- Carter, E., Schönbrodt, F., Gervais, W. M., & Hilgard, J. (2017). Correcting for bias in psychology: A comparison of meta-analytic methods.
- Chiang, H.-T. L., Faust, A., & Tapia, L. (2018). Deep neural networks for swept volume prediction between configurations. *arXiv preprint arXiv:1805.11597*.
- Chollet, F. (2017). *Deep learning with python*. Manning Publications Co.
- Chollet, F. et al. (2018). Keras: The python deep learning library. *Astrophysics Source Code Library*.
- Coburn, K. M. & Vevea, J. L. (2017). *Weighttr: Estimating weight-function models for publication bias*. R package version 1.1.2. Retrieved from <https://CRAN.R-project.org/package=weighttr>

- Cohen, J. (1990). Things i have learned (so far). *American psychologist*, 45(12), 1304.
- Copas, J. (1999). What works?: Selectivity models and meta-analysis. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 162(1), 95–109.
- Copas, J. & Shi, J. Q. (2000). Meta-analysis, funnel plots and sensitivity analysis. *Biostatistics*, 1(3), 247–262.
- de Vrieze, J. (2018). The metawars. *Science*.
- Du, H., Liu, F., & Wang, L. (2017). A bayesian “fill-in” method for correcting for publication bias in meta-analysis. *Psychological methods*, 22(4), 799.
- Duval, S. & Tweedie, R. (2000). Trim and fill: A simple funnel-plot-based method of testing and adjusting for publication bias in meta-analysis. *Biometrics*, 56(2), 455–463.
- Dwan, K., Altman, D. G., Arnaiz, J. A., Bloom, J., Chan, A.-W., Cronin, E., . . . Gamble, C., et al. (2008). Systematic review of the empirical evidence of study publication bias and outcome reporting bias. *PloS one*, 3(8), e3081.
- Fanelli, D., Costas, R., & Ioannidis, J. P. (2017). Meta-assessment of bias in science. *Proceedings of the National Academy of Sciences*, 201618569.
- Franco, A., Malhotra, N., & Simonovits, G. (2014). Publication bias in the social sciences: Unlocking the file drawer. *Science*, 345(6203), 1502–1505.
- Fritz, C. O., Morris, P. E., & Richler, J. J. (2012). Effect size estimates: Current use, calculations, and interpretation. *Journal of experimental psychology: General*, 141(1), 2.
- Gal, Y. & Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International conference on machine learning* (pp. 1050–1059).
- Hedges, L. V. & Vevea, J. L. (1998). Fixed-and random-effects models in meta-analysis. *Psychological methods*, 3(4), 486.
- Ioannidis, J. P. & Trikalinos, T. A. (2007). An exploratory test for an excess of significant findings. *Clinical trials*, 4(3), 245–253.

- Ioannidis, J., Stanley, T., & Doucouliagos, H. (2017). The power of bias in economics research. *The Economic Journal*, 127(605).
- Ioffe, S. & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- Iyengar, S. & Greenhouse, J. B. (1988). Selection models and the file drawer problem. *Statistical Science*, 109–117.
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 0956797611430953.
- Kahneman, D. (2003). Experiences of collaborative research. *American Psychologist*, 58(9), 723.
- Kuhn, M. (2008). Building predictive models in r using the caret package. *Journal of Statistical Software, Articles*, 28(5), 1–26. doi:[10.18637/jss.v028.i05](https://doi.org/10.18637/jss.v028.i05)
- Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems* (pp. 6402–6413).
- Marszalek, J. M., Barber, C., Kohlhart, J., & Cooper, B. H. (2011). Sample size in psychological research over the past 30 years. *Perceptual and motor skills*, 112(2), 331–348.
- McShane, B. B., Böckenholt, U., & Hansen, K. T. (2016). Adjusting for publication bias in meta-analysis: An evaluation of selection methods and some cautionary notes. *Perspectives on Psychological Science*, 11(5), 730–749.
- Moreno, S. G., Sutton, A. J., Ades, A., Stanley, T., Abrams, K. R., Peters, J. L., & Cooper, N. J. (2009). Assessment of regression-based methods to adjust for publication bias through a comprehensive simulation study. *BMC Medical Research Methodology*, 9(1), 2.
- Pan, S. J., Yang, Q. et al. (2010). A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10), 1345–1359.

- Pearce, T., Zaki, M., Brintrup, A., & Neel, A. (2018). Uncertainty in neural networks: Bayesian ensembling. *arXiv preprint arXiv:1810.05546*.
- R Core Team. (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological bulletin*, 86(3), 638.
- Rücker, G., Schwarzer, G., Carpenter, J. R., Binder, H., & Schumacher, M. (2010). Treatment-effect estimates adjusted for small-study effects via a limit meta-analysis. *Biostatistics*, 12(1), 122–142.
- Simons, D. J., Holcombe, A. O., & Spellman, B. A. (2014). An introduction to registered replication reports at perspectives on psychological science. *Perspectives on Psychological Science*, 9(5), 552–555.
- Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve and effect size: Correcting for publication bias using only significant results. *Perspectives on Psychological Science*, 9(6), 666–681.
- Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2015). Better p-curves: Making p-curve analysis more robust to errors, fraud, and ambitious p-hacking, a reply to ulrich and miller (2015).
- Stanley, T. & Doucouliagos, H. (2014). Meta-regression approximations to reduce publication selection bias. *Research Synthesis Methods*, 5(1), 60–78.
- Stanley, T., Doucouliagos, H., & Ioannidis, J. (2017). Finding the power to reduce publication bias. *Statistics in Medicine*, 36(10), 1580–1598.
- Stegenga, J. (2011). Is meta-analysis the platinum standard of evidence? *Studies in history and philosophy of science part C: Studies in history and philosophy of biological and biomedical sciences*, 42(4), 497–507.
- Sterne, J. A., Gavaghan, D., & Egger, M. (2000). Publication and related bias in meta-analysis: Power of statistical tests and prevalence in the literature. *Journal of clinical Epidemiology*, 53(11), 1119–1129.

- Tran, D., Kucukelbir, A., Dieng, A. B., Rudolph, M., Liang, D., & Blei, D. M. (2016). Edward: A library for probabilistic modeling, inference, and criticism. *arXiv preprint arXiv:1610.09787*.
- van Aert, R. C. (2017). *Puniform: Meta-analysis methods correcting for publication bias*. R package version 0.0.3.
- van Aert, R. & van Assen, M. (2018). Correcting for publication bias in a meta-analysis with the p-uniform* method. *OSF preprint* <https://osf.io/jdb6k/>.
- van Aert, R. C., Wicherts, J. M., & van Assen, M. A. (2016). Conducting meta-analyses based on p values: Reservations and recommendations for applying p-uniform and p-curve. *Perspectives on Psychological Science*, 11(5), 713–729.
- Viechtbauer, W. (2005). Bias and efficiency of meta-analytic variance estimators in the random-effects model. *Journal of Educational and Behavioral Statistics*, 30(3), 261–293.
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(3), 1–48.

Figure 1. The six candidate network architectures for DeepMA. The input layer always included 118 units (one unit per feature) and fed into a batch normalization layer (Ioffe & Szegedy, 2015). Layers labeled 1“Full”are standard fully connected layers. The final layer included two units, one for $\hat{\delta}$ and one for $\hat{\tau}$

MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank		
0.05	D1	1	0.02	D1	1	0.02	D1	1	0.08	D1	1	0.06	D1	1
0.12	PE	2	0.06	PP	2	0.03	RE	2	0.23	PE	2	0.07	RE	2
0.12	WW	3	0.07	PE	3	0.03	TF	3	0.23	PT	2	0.1	TF	3
0.14	PP	4	0.08	PU	4	0.03	WW	3	0.23	PP	2	0.11	TP	4
0.14	TF	5	0.08	FP	4	0.03	TP	4	0.26	WW	3	0.14	WW	5
0.15	PT	6	0.09	TP	4	0.04	PE	5	0.31	TF	4	0.15	FP	5
0.16	TP	6	0.09	WW	5	0.04	FP	6	0.33	FP	5	0.17	PE	6
0.17	FP	7	0.09	PC	5	0.05	PP	7	0.37	TP	6	0.22	PP	7
0.19	RE	8	0.1	PT	6	0.06	PT	8	0.41	PU	6	0.25	PT	8
0.25	PU	8	0.14	TF	7	0.1	PU	9	0.42	RE	7	0.53	PU	9
0.31	PC	9	0.19	RE	8	0.21	PC	10	0.44	PC	7	0.58	PC	9
0.16	D1	1	0.03	FP	1	0.03	WW	1	0.31	D1	1	0.15	D1	1
0.2	WW	2	0.03	TP	1	0.03	RE	1	0.32	PU	1	0.15	RE	1
0.21	PE	3	0.07	D1	2	0.03	TF	1	0.33	PT	1	0.16	TF	1
0.24	TF	3	0.08	PE	3	0.04	PE	2	0.33	PC	1	0.19	WW	2
0.26	PP	4	0.08	PP	3	0.04	TP	2	0.36	PP	1	0.22	PE	3
0.27	PU	5	0.08	PU	4	0.05	D1	2	0.37	PE	2	0.28	PP	4
0.27	RE	5	0.08	WW	5	0.06	FP	3	0.38	WW	3	0.31	PT	5
0.27	PT	6	0.09	TF	6	0.06	PP	3	0.45	TF	4	0.43	TP	6
0.34	FP	6	0.11	PC	6	0.06	PU	4	0.53	TP	4	0.45	FP	6
0.34	TP	6	0.12	RE	7	0.07	PT	4	0.53	FP	4	0.65	PU	7
0.5	PC	6	0.13	PT	8	1	PC	5	0.56	RE	5	0.83	PC	8
All			$\tau \leq 0.10$, High bias			$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias		

Figure 2. Mean Absolute Error (MAE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets). D1 = the initial version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank		
0.05	D2	1	0.02	D2	1	0.02	D2	1	0.09	D2	1	0.06	D2	1
0.12	PE	2	0.06	PP	2	0.03	RE	2	0.23	PE	2	0.07	RE	2
0.12	WW	3	0.07	PE	3	0.03	TF	3	0.23	PT	2	0.1	TF	3
0.14	PP	4	0.08	PU	4	0.03	WW	3	0.23	PP	2	0.11	TP	4
0.14	TF	5	0.08	FP	4	0.03	TP	4	0.26	WW	3	0.14	WW	5
0.15	PT	6	0.09	TP	4	0.04	PE	5	0.31	TF	4	0.15	FP	5
0.16	TP	6	0.09	WW	5	0.04	FP	6	0.33	FP	5	0.17	PE	6
0.17	FP	7	0.09	PC	5	0.05	PP	7	0.37	TP	6	0.22	PP	7
0.19	RE	8	0.1	PT	6	0.06	PT	8	0.41	PU	6	0.25	PT	8
0.25	PU	8	0.14	TF	7	0.1	PU	9	0.42	RE	7	0.53	PU	9
0.31	PC	9	0.19	RE	8	0.21	PC	10	0.44	PC	7	0.58	PC	9
Initial test set														
0.03	D2	1	0	D2	1	0.01	D2	1	0.06	D2	1	0.05	D2	1
0.2	WW	2	0.03	FP	2	0.03	WW	2	0.32	PU	2	0.15	RE	2
0.21	PE	3	0.03	TP	2	0.03	RE	2	0.33	PT	2	0.16	TF	3
0.24	TF	3	0.08	PE	3	0.03	TF	2	0.33	PC	2	0.19	WW	4
0.26	PP	4	0.08	PP	3	0.04	PE	3	0.36	PP	3	0.22	PE	5
0.27	PU	5	0.08	PU	4	0.04	TP	3	0.37	PE	4	0.28	PP	6
0.27	RE	5	0.08	WW	5	0.06	FP	4	0.38	WW	5	0.31	PT	7
0.27	PT	6	0.09	TF	6	0.06	PP	4	0.45	TF	6	0.43	TP	8
0.34	FP	6	0.11	PC	6	0.06	PU	5	0.53	TP	6	0.45	FP	8
0.34	TP	6	0.12	RE	7	0.07	PT	5	0.53	FP	6	0.65	PU	9
0.5	PC	6	0.13	PT	8	1	PC	6	0.56	RE	7	0.83	PC	10
Transfer test set														
All			$\tau \leq 0.10$, High bias			$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias		

Figure 3. Mean Absolute Error (MAE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets). D2 = the re-trained version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank		
0.04	D1	1	0.02	D1	1	0.03	D1	1	0.06	D1	1	0.04	D1	1
0.07	TP	2	0.03	TP	1	0.03	TP	1	0.12	TF	2	0.06	RE	2
0.07	FP	3	0.03	FP	1	0.03	FP	1	0.12	FP	2	0.07	TP	3
0.07	RE	4	0.06	TF	2	0.03	RE	2	0.13	TP	3	0.08	TF	4
0.08	TF	5	0.07	RE	3	0.04	TF	3	0.16	RE	4	0.08	FP	4
Initial test set														
0.21	TF	1	0.01	TP	1	0.02	TP	1	0.35	TF	1	0.28	TP	1
0.21	RE	1	0.01	FP	1	0.03	FP	1	0.4	D1	2	0.28	RE	1
0.22	D1	2	0.03	RE	1	0.03	RE	1	0.4	RE	3	0.29	TF	1
0.25	TP	3	0.04	D1	1	0.04	TF	1	0.41	TP	3	0.31	FP	2
0.25	FP	3	0.04	TF	1	0.04	D1	1	0.41	FP	3	0.34	D1	3
Transfer test set														
All			$\tau \leq 0.10$, High bias			$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias		

Figure 4. Mean Absolute Error (MAE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets). D2 = the re-trained version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank		
0.04	D2	1	0.03	TP	1	0.03	TP	1	0.06	D2	1	0.05	D2	1
0.07	TP	2	0.03	FP	1	0.03	FP	1	0.12	TF	2	0.06	RE	2
0.07	FP	3	0.03	D2	2	0.03	RE	2	0.12	FP	2	0.07	TP	3
0.07	RE	4	0.06	TF	3	0.03	D2	3	0.13	TP	3	0.08	TF	4
0.08	TF	5	0.07	RE	4	0.04	TF	3	0.16	RE	4	0.08	FP	4
Initial test set														
0.04	D2	1	0.01	TP	1	0.02	TP	1	0.1	D2	1	0.07	D2	1
0.21	TF	2	0.01	FP	1	0.03	FP	1	0.35	TF	2	0.28	TP	2
0.21	RE	2	0.02	D2	1	0.03	D2	1	0.4	RE	3	0.28	RE	2
0.25	TP	3	0.03	RE	1	0.03	RE	1	0.41	TP	3	0.29	TF	2
0.25	FP	3	0.04	TF	1	0.04	TF	1	0.41	FP	3	0.31	FP	3
Transfer test set														
All			$\tau \leq 0.10$, High bias			$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias		

Figure 5. Mean Absolute Error (MAE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets). D2 = the re-trained version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			
0.07	D1	1	0.04	D1	1	0.03	D1	1	0.13	D1	1	0.09	D1	1	Initial test set
0.14	TF	2	0.08	TP	2	0.05	RE	2	0.19	PU	2	0.12	RE	2	
0.15	WW	3	0.08	FP	2	0.05	TF	2	0.19	PC	2	0.13	TF	2	
0.15	TP	3	0.09	PE	2	0.06	WW	2	0.22	PE	3	0.19	WW	3	
0.15	PE	3	0.11	PC	3	0.07	TP	3	0.24	TF	3	0.21	TP	3	
0.16	FP	4	0.11	PU	3	0.08	PE	4	0.24	WW	3	0.25	PE	4	
0.17	PU	4	0.12	PP	3	0.09	FP	4	0.24	TP	3	0.28	FP	4	
0.18	RE	5	0.14	TF	4	0.12	PP	5	0.25	FP	3	0.35	PP	5	
0.22	PC	5	0.14	WW	4	0.13	PU	5	0.29	PP	4	0.36	PU	5	
0.22	PP	6	0.15	PT	5	0.14	PT	6	0.3	PT	4	0.4	PT	6	
0.24	PT	7	0.19	RE	6	0.47	PC	6	0.32	RE	5	0.48	PC	6	
0.17	D1	1	0.05	FP	1	0.05	WW	1	0.3	D1	1	0.2	D1	1	Transfer test set
0.23	WW	2	0.05	TP	1	0.05	RE	1	0.34	PU	2	0.26	RE	2	
0.26	PE	2	0.07	D1	1	0.06	TF	2	0.36	PC	2	0.26	TF	2	
0.27	TF	3	0.08	PE	2	0.07	D1	2	0.4	PE	3	0.29	WW	3	
0.3	RE	4	0.09	PP	2	0.07	PE	3	0.41	PT	3	0.35	PE	4	
0.3	PU	4	0.09	PU	3	0.08	TP	3	0.42	PP	3	0.46	PP	5	
0.32	TP	4	0.11	WW	4	0.09	PU	4	0.42	WW	4	0.48	PT	5	
0.33	FP	4	0.11	TF	4	0.1	FP	4	0.45	TP	4	0.49	TP	5	
0.34	PP	5	0.13	PT	5	0.1	PP	4	0.45	FP	4	0.52	FP	6	
0.36	PT	6	0.14	PC	5	0.12	PT	5	0.47	TF	5	0.73	PU	7	
0.76	PC	7	0.15	RE	6	2.06	PC	6	0.55	RE	6	1.1	PC	8	
All			$\tau \leq 0.10$, High bias			$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias			

Figure 6. Mean Absolute Error (MAE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets) when $k \leq 10$. D1 = the initial version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			
0.08	D2	1	0.04	D2	1	0.03	D2	1	0.14	D2	1	0.1	D2	1	Initial test set
0.14	TF	2	0.08	TP	2	0.05	RE	2	0.19	PU	2	0.12	RE	2	
0.15	WW	3	0.08	FP	2	0.05	TF	2	0.19	PC	2	0.13	TF	2	
0.15	TP	3	0.09	PE	2	0.06	WW	2	0.22	PE	3	0.19	WW	3	
0.15	PE	3	0.11	PC	3	0.07	TP	3	0.24	TF	3	0.21	TP	3	
0.16	FP	4	0.11	PU	3	0.08	PE	4	0.24	WW	3	0.25	PE	4	
0.17	PU	4	0.12	PP	3	0.09	FP	4	0.24	TP	3	0.28	FP	4	
0.18	RE	5	0.14	TF	4	0.12	PP	5	0.25	FP	3	0.35	PP	5	
0.22	PC	5	0.14	WW	4	0.13	PU	5	0.29	PP	4	0.36	PU	5	
0.22	PP	6	0.15	PT	5	0.14	PT	6	0.3	PT	4	0.4	PT	6	
0.24	PT	7	0.19	RE	6	0.47	PC	6	0.32	RE	5	0.48	PC	6	
0.06	D2	1	0.01	D2	1	0.01	D2	1	0.12	D2	1	0.1	D2	1	Transfer test set
0.23	WW	2	0.05	FP	2	0.05	WW	2	0.34	PU	2	0.26	RE	2	
0.26	PE	2	0.05	TP	2	0.05	RE	2	0.36	PC	2	0.26	TF	2	
0.27	TF	3	0.08	PE	3	0.06	TF	3	0.4	PE	3	0.29	WW	3	
0.3	RE	4	0.09	PP	3	0.07	PE	4	0.41	PT	3	0.35	PE	4	
0.3	PU	4	0.09	PU	4	0.08	TP	4	0.42	PP	3	0.46	PP	5	
0.32	TP	4	0.11	WW	5	0.09	PU	5	0.42	WW	4	0.48	PT	5	
0.33	FP	4	0.11	TF	5	0.1	FP	5	0.45	TP	4	0.49	TP	5	
0.34	PP	5	0.13	PT	6	0.1	PP	5	0.45	FP	4	0.52	FP	6	
0.36	PT	6	0.14	PC	6	0.12	PT	6	0.47	TF	5	0.73	PU	7	
0.76	PC	7	0.15	RE	7	2.06	PC	7	0.55	RE	6	1.1	PC	8	
All			$\tau \leq 0.10$, High bias			$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias			

Figure 7. Mean Absolute Error (MAE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets) when $k \leq 10$. D2 = the re-trained version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank		
0.05	D1	1	0.02	D1	1	0.02	D1	1	0.09	D1	1	0.05	D1	1
0.14	PE	2	0.07	PP	2	0.03	RE	2	0.24	PT	2	0.07	RE	2
0.15	PT	2	0.08	PU	2	0.03	WW	2	0.27	PP	3	0.09	TF	3
0.15	PP	3	0.08	PT	3	0.03	TF	3	0.28	PE	4	0.11	TP	4
0.15	WW	4	0.1	PE	4	0.03	PE	4	0.33	FP	5	0.14	WW	5
0.18	TP	5	0.11	PC	4	0.03	TP	4	0.33	WW	6	0.16	FP	6
0.19	FP	6	0.12	FP	4	0.05	FP	5	0.37	TP	6	0.16	PE	7
0.19	TF	7	0.12	TP	4	0.05	PP	5	0.4	TF	7	0.22	PP	8
0.25	RE	8	0.14	WW	5	0.05	PT	6	0.45	PU	7	0.25	PT	9
0.28	PU	8	0.24	TF	6	0.15	PU	7	0.48	PC	7	0.57	PU	10
0.36	PC	8	0.32	RE	7	0.36	PC	8	0.53	RE	8	0.62	PC	10
Initial test set														
0.16	D1	1	0.02	FP	1	0.03	WW	1	0.33	PU	1	0.1	D1	1
0.23	PE	2	0.02	TP	1	0.03	RE	1	0.33	PT	1	0.14	RE	1
0.23	WW	3	0.09	PT	2	0.03	TF	1	0.34	D1	2	0.14	TF	1
0.24	PT	3	0.1	PU	3	0.04	PE	2	0.35	PC	2	0.15	WW	1
0.26	PP	4	0.1	D1	4	0.05	TP	3	0.37	PP	3	0.18	PE	2
0.3	TF	5	0.1	PP	4	0.06	D1	4	0.42	PE	4	0.24	PP	3
0.31	PU	5	0.11	PE	5	0.06	PP	4	0.45	WW	5	0.25	PT	3
0.33	RE	6	0.13	PC	5	0.07	PT	4	0.51	TP	5	0.51	TP	4
0.41	FP	6	0.14	WW	6	0.08	FP	4	0.51	FP	5	0.54	FP	4
0.41	TP	6	0.16	TF	7	0.09	PU	5	0.54	TF	6	0.81	PU	5
0.59	PC	7	0.19	RE	8	1.22	PC	6	0.65	RE	7	1.01	PC	6
Transfer test set														
All			$\tau \leq 0.10$, High bias			$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias		

Figure 8. Mean Absolute Error (MAE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets) when $\delta \leq 0.20$. D1 = the initial version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank			MAE Method Rank		
0.05	D2	1	0.03	D2	1	0.02	D2	1	0.1	D2	1	0.05	D2	1
0.14	PE	2	0.07	PP	2	0.03	RE	2	0.24	PT	2	0.07	RE	2
0.15	PT	2	0.08	PU	2	0.03	WW	2	0.27	PP	3	0.09	TF	3
0.15	PP	3	0.08	PT	3	0.03	TF	3	0.28	PE	4	0.11	TP	4
0.15	WW	4	0.1	PE	4	0.03	PE	4	0.33	FP	5	0.14	WW	5
0.18	TP	5	0.11	PC	4	0.03	TP	4	0.33	WW	6	0.16	FP	6
0.19	FP	6	0.12	FP	4	0.05	FP	5	0.37	TP	6	0.16	PE	7
0.19	TF	7	0.12	TP	4	0.05	PP	5	0.4	TF	7	0.22	PP	8
0.25	RE	8	0.14	WW	5	0.05	PT	6	0.45	PU	7	0.25	PT	9
0.28	PU	8	0.24	TF	6	0.15	PU	7	0.48	PC	7	0.57	PU	10
0.36	PC	8	0.32	RE	7	0.36	PC	8	0.53	RE	8	0.62	PC	10
Initial test set														
0.03	D2	1	0	D2	1	0.01	D2	1	0.07	D2	1	0.04	D2	1
0.23	PE	2	0.02	FP	2	0.03	WW	2	0.33	PU	2	0.14	RE	2
0.23	WW	3	0.02	TP	2	0.03	RE	2	0.33	PT	2	0.14	TF	2
0.24	PT	3	0.09	PT	3	0.03	TF	2	0.35	PC	2	0.15	WW	3
0.26	PP	4	0.1	PU	4	0.04	PE	3	0.37	PP	3	0.18	PE	4
0.3	TF	5	0.1	PP	5	0.05	TP	4	0.42	PE	4	0.24	PP	5
0.31	PU	5	0.11	PE	6	0.06	PP	4	0.45	WW	5	0.25	PT	5
0.33	RE	6	0.13	PC	6	0.07	PT	5	0.51	TP	5	0.51	TP	6
0.41	FP	6	0.14	WW	7	0.08	FP	5	0.51	FP	5	0.54	FP	6
0.41	TP	6	0.16	TF	8	0.09	PU	6	0.54	TF	6	0.81	PU	7
0.59	PC	7	0.19	RE	9	1.22	PC	7	0.65	RE	7	1.01	PC	8
Transfer test set														
All			$\tau \leq 0.10$, High bias			$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias		

Figure 9. Mean Absolute Error (MAE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets) when $\delta \leq 0.20$. D2 = the re-trained version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

Figure 10. Performance in terms of Mean Absolute Error (MAE) for the initial (plus signs) and re-trained (squares) versions of DeepMA. The columns are either all of the test data or subsets of those data representing the combinations of low or high heterogeneity and none or high bias. The rows are all data and subsets of either small values for k and δ . The horizontal axis for each facet represents either the initial or transfer test set. The dashed line represents the upper limit of the good estimation performance range based on the performance of the random-effects meta-analysis model on the initial test set and in the absence of bias.

ME	Method	Rank	ME	Method	Rank	ME	Method	Rank	ME	Method	Rank	ME	Method	Rank
0	PP	1	0	D1	1	0	RE	1	0	PT	1	0	RE	1
0	D1	1	-0.01	PP	1	0	TP	1	0.01	D1	1	0	D1	1
0.01	TF	2	0.02	PE	2	0	WW	1	0.04	PP	2	0	WW	1
0.01	PE	2	0.02	WW	3	0	D1	2	0.04	TF	3	-0.01	TP	1
0.02	WW	3	0.04	TF	4	0	FP	2	0.07	PE	3	-0.01	FP	1
-0.03	PT	3	-0.05	PC	4	0	PE	2	0.07	WW	3	-0.01	PE	1
0.08	RE	4	-0.05	FP	4	-0.01	PT	2	0.18	RE	4	-0.02	PT	1
-0.12	TP	4	-0.05	TP	4	-0.01	PP	2	-0.29	FP	4	-0.05	PP	1
-0.13	FP	4	-0.06	PT	4	-0.01	TF	2	-0.33	TP	4	-0.05	TF	1
0.2	PU	5	-0.06	PU	4	0.06	PU	3	0.38	PU	5	0.52	PU	2
0.26	PC	6	0.11	RE	5	0.17	PC	4	0.42	PC	6	0.57	PC	2

0.04	D1	1	-0.01	FP	1	0	WW	1	0.14	PT	1	-0.01	RE	1
-0.05	PP	1	-0.01	TP	1	0	RE	1	0.17	D1	2	-0.04	TF	1
0.07	PE	2	-0.03	PP	1	0	PE	1	0.18	PU	2	-0.06	WW	1
-0.1	PT	2	0.04	D1	2	0	D1	2	0.21	PC	3	-0.06	D1	1
-0.1	FP	2	0.05	PE	2	-0.01	TF	2	0.21	PP	4	-0.07	PE	1
-0.11	TP	2	-0.05	PC	2	-0.01	PT	2	-0.22	TP	4	-0.13	PP	1
0.11	WW	3	-0.07	PU	2	0.01	TP	2	-0.22	FP	4	-0.13	PT	1
0.18	TF	4	0.07	WW	3	0.01	FP	2	0.33	PE	5	0.22	TP	2
0.18	PU	4	0.08	TF	4	0.01	PU	2	0.36	WW	6	0.24	FP	3
0.23	RE	5	0.12	RE	5	-0.01	PP	2	0.44	TF	7	0.62	PU	4
0.42	PC	5	-0.12	PT	5	0.97	PC	3	0.55	RE	8	0.81	PC	5

All	$\tau \leq 0.10$, High bias	$\tau \leq 0.10$, No bias	$\tau \geq 0.40$, High bias	$\tau \geq 0.40$, No bias

Figure 11. Mean error (ME) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets). D1 = the initial version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

ME	Method	Rank	ME	Method	Rank	ME	Method	Rank	ME	Method	Rank	ME	Method	Rank
0	PP	1	-0.01	D2	1	0	RE	1	0	PT	1	0	RE	1
0	D2	1	-0.01	PP	1	0	TP	1	0.02	D2	2	0	WW	1
0.01	TF	2	0.02	PE	2	0	WW	1	0.04	PP	3	-0.01	D2	1
0.01	PE	2	0.02	WW	3	0	FP	1	0.04	TF	4	-0.01	TP	1
0.02	WW	3	0.04	TF	4	0	PE	1	0.07	PE	4	-0.01	FP	1
-0.03	PT	3	-0.05	PC	4	0	D2	1	0.07	WW	4	-0.01	PE	1
0.08	RE	4	-0.05	FP	4	-0.01	PT	1	0.18	RE	5	-0.02	PT	1
-0.12	TP	4	-0.05	TP	4	-0.01	PP	1	-0.29	FP	5	-0.05	PP	1
-0.13	FP	4	-0.06	PT	4	-0.01	TF	1	-0.33	TP	5	-0.05	TF	1
0.2	PU	5	-0.06	PU	4	0.06	PU	2	0.38	PU	6	0.52	PU	2
0.26	PC	6	0.11	RE	5	0.17	PC	3	0.42	PC	7	0.57	PC	2

0	D2	1	0	D2	1	0	WW	1	0.01	D2	1	-0.01	RE	1
-0.05	PP	1	-0.01	FP	1	0	RE	1	0.14	PT	2	-0.01	D2	2
0.07	PE	2	-0.01	TP	1	0	D2	2	0.18	PU	3	-0.04	TF	2
-0.1	PT	2	-0.03	PP	1	0	PE	2	0.21	PC	4	-0.06	WW	2
-0.1	FP	2	0.05	PE	2	-0.01	TF	2	0.21	PP	5	-0.07	PE	2
-0.11	TP	2	-0.05	PC	2	-0.01	PT	2	-0.22	TP	5	-0.13	PP	2
0.11	WW	3	-0.07	PU	2	0.01	TP	2	-0.22	FP	5	-0.13	PT	2
0.18	TF	4	0.07	WW	3	0.01	FP	2	0.33	PE	6	0.22	TP	3
0.18	PU	4	0.08	TF	4	0.01	PU	2	0.36	WW	7	0.24	FP	4
0.23	RE	5	0.12	RE	5	-0.01	PP	2	0.44	TF	8	0.62	PU	5
0.42	PC	5	-0.12	PT	5	0.97	PC	3	0.55	RE	9	0.81	PC	6

All	$\tau \leq 0.10$, High bias	$\tau \leq 0.10$, No bias	$\tau \geq 0.40$, High bias	$\tau \geq 0.40$, No bias

Figure 12. Mean error (ME) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets). D2 = the re-trained version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

Method Rank RMSE			Method Rank RMSE			Method Rank RMSE			Method Rank RMSE			Method Rank RMSE			Initial test set
D1	1	0.09	D1	1	0.05	RE	1	0.04	D1	1	0.15	RE	1	0.1	
PE	2	0.17	PP	2	0.09	TF	2	0.04	PE	2	0.29	D1	1	0.1	
WW	3	0.18	PE	3	0.1	WW	2	0.04	PP	2	0.31	TF	2	0.14	
TF	4	0.21	PU	4	0.12	D1	2	0.04	PT	2	0.32	TP	3	0.19	
PP	4	0.21	PT	5	0.12	TP	3	0.05	WW	3	0.32	WW	4	0.2	
PT	5	0.23	FP	5	0.12	PE	4	0.06	TF	4	0.37	FP	4	0.23	
FP	6	0.25	TP	5	0.13	FP	5	0.06	FP	5	0.42	PE	5	0.23	
TP	6	0.26	WW	5	0.13	PP	6	0.08	RE	6	0.47	PP	6	0.31	
RE	7	0.26	TF	6	0.19	PT	7	0.1	TP	6	0.47	PT	7	0.35	
PU	7	0.42	RE	7	0.25	PU	8	0.23	PU	6	0.61	PU	8	0.67	
PC	8	0.66	PC	7	0.28	PC	9	0.67	PC	6	0.69	PC	8	0.82	
D1	1	0.24	FP	1	0.05	WW	1	0.05	D1	1	0.38	D1	1	0.2	Transfer test set
WW	2	0.32	TP	1	0.05	RE	1	0.06	PU	1	0.4	RE	1	0.26	
PE	3	0.37	D1	2	0.09	TF	1	0.06	PC	1	0.43	TF	1	0.26	
TF	3	0.45	PE	3	0.09	D1	2	0.08	WW	2	0.46	WW	2	0.28	
PU	4	0.47	PU	4	0.11	TP	2	0.1	PE	2	0.47	PE	3	0.4	
RE	4	0.48	PP	4	0.11	PE	2	0.11	PT	2	0.5	PP	4	0.61	
PP	4	0.53	WW	5	0.11	PU	3	0.13	PP	2	0.53	PT	5	0.65	
PT	5	0.54	TF	6	0.12	PP	3	0.22	TF	3	0.53	PU	6	0.92	
TP	5	0.81	PT	7	0.15	PT	3	0.22	RE	4	0.62	TP	6	0.93	
FP	5	0.82	RE	7	0.16	FP	3	0.23	TP	4	0.76	FP	6	0.95	
PC	5	1.03	PC	7	0.32	PC	4	1.88	FP	4	0.76	PC	7	1.17	
All	$\tau \leq 0.10$, High bias					$\tau \leq 0.10$, No bias			$\tau \geq 0.40$, High bias			$\tau \geq 0.40$, No bias			

Figure 13. Root Mean Squared Error (RMSE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets). D1 = the initial version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.

Method Rank RMSE	Method Rank RMSE	Method Rank RMSE	Method Rank RMSE	Method Rank RMSE	Initial test set
D2 1 0.09	D2 1 0.05	RE 1 0.04	D2 1 0.15	RE 1 0.1	
PE 2 0.17	PP 2 0.09	TF 2 0.04	PE 2 0.29	D2 1 0.11	
WW 3 0.18	PE 3 0.1	WW 2 0.04	PP 2 0.31	TF 2 0.14	
TF 4 0.21	PU 4 0.12	D2 2 0.04	PT 2 0.32	TP 3 0.19	
PP 4 0.21	PT 5 0.12	TP 3 0.05	WW 3 0.32	WW 4 0.2	
PT 5 0.23	FP 5 0.12	PE 4 0.06	TF 4 0.37	FP 4 0.23	
FP 6 0.25	TP 5 0.13	FP 5 0.06	FP 5 0.42	PE 5 0.23	
TP 6 0.26	WW 5 0.13	PP 6 0.08	RE 6 0.47	PP 6 0.31	
RE 7 0.26	TF 6 0.19	PT 7 0.1	TP 6 0.47	PT 7 0.35	
PU 7 0.42	RE 7 0.25	PU 8 0.23	PU 6 0.61	PU 8 0.67	
PC 8 0.66	PC 7 0.28	PC 9 0.67	PC 6 0.69	PC 8 0.82	
D2 1 0.1	D2 1 0.03	D2 1 0.04	D2 1 0.15	D2 1 0.13	Transfer test set
WW 2 0.32	FP 2 0.05	WW 2 0.05	PU 2 0.4	RE 2 0.26	
PE 3 0.37	TP 2 0.05	RE 2 0.06	PC 2 0.43	TF 3 0.26	
TF 3 0.45	PE 3 0.09	TF 2 0.06	WW 3 0.46	WW 4 0.28	
PU 4 0.47	PU 4 0.11	TP 3 0.1	PE 3 0.47	PE 5 0.4	
RE 4 0.48	PP 4 0.11	PE 3 0.11	PT 3 0.5	PP 6 0.61	
PP 4 0.53	WW 5 0.11	PU 4 0.13	PP 3 0.53	PT 7 0.65	
PT 5 0.54	TF 6 0.12	PP 4 0.22	TF 4 0.53	PU 8 0.92	
TP 5 0.81	PT 7 0.15	PT 4 0.22	RE 5 0.62	TP 8 0.93	
FP 5 0.82	RE 7 0.16	FP 4 0.23	TP 5 0.76	FP 8 0.95	
PC 5 1.03	PC 7 0.32	PC 5 1.88	FP 5 0.76	PC 9 1.17	
All	$\tau \leq 0.10$, High bias	$\tau \leq 0.10$, No bias	$\tau \geq 0.40$, High bias	$\tau \geq 0.40$, No bias	

Figure 14. Root Mean Squared Error (RMSE) for all methods across two test sets (shown in terms of either all of the data in the set or as four subsets). D2 = the re-trained version of DeepMA; FP = four parameter selection model; TP = three parameter selection model; TF = trim-and-fill; RE = random-effects meta-analysis; PE = precision effect estimate with standard error (PEESE); PT = precision effect test (PET); PP = PET-PEESE; PU = p -uniform; PC = p -curve; WW = weighted average of the adequately powered.