Millions of Findings at Your Fingertips

February 26, 2020

met•a•bus (ˈme-tə-ˌbəs) *n.*
1. A cloud-based research synthesis platform sitting atop the world's largest collection of curated social science research findings.

1,113,483 Findings   14,081 Articles   1,175 Users
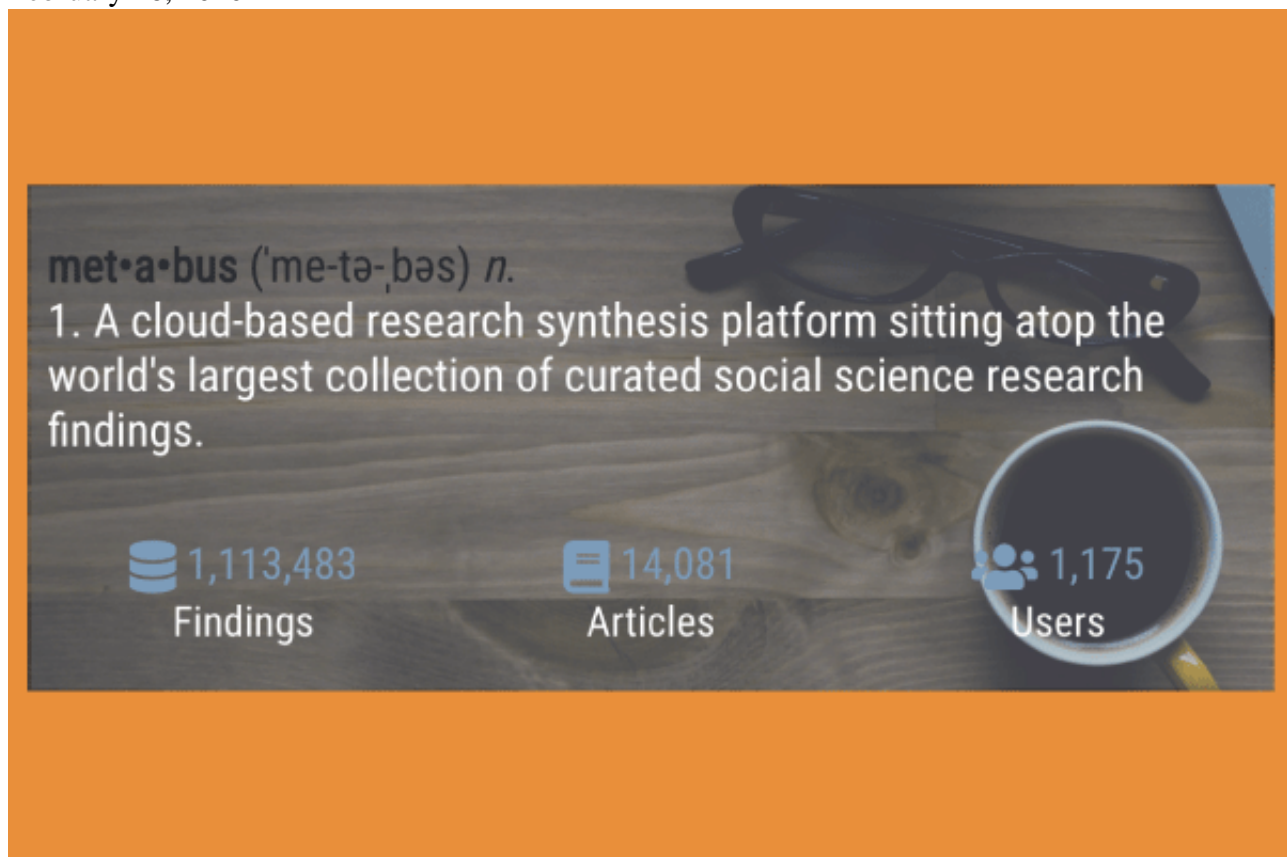
Thefield of psychology has produced hundreds of thousands of journal articles and,within them, millions of findings in the form of effect sizes (e.g.,correlation coefficients). As researchers, we rely on these findingspractically every day—to estimate statistical power, generate or justifyresearch questions, conduct meta-analyses, interpret research outcomes, satisfysimple curiosities, and, for some, to joke with colleagues about theconventions of null-hypothesis significance testing. After all, as noted byCohen (1988), the effect size is "what science is all about" (p. 532). So, inan age in which most of us are accustomed to Amazon.com, why do we still searchfor research findings the old-fashioned way? Why don't we have a search engineof our research findings?

Most psychologists are familiar withthe traditional, "top-down" literature search process. It starts by specifyinga phenomenon of interest and ends after arduous hours of searching andfiltering. Although impressive in scope, search engines like Google Scholaroperate at a higher order (i.e., article) level and, consequently, often returntoo few hits and too many misses. This can make the search process a monumentof inefficiency.

But what if things were different?What if the millions of psychology findings were managed with Amazon.com-likeefficiency? Imagine if a researcher could visit a website, request all reportedfindings on the relation between, say, conscientiousness and age, and then viewan instant meta-analysis

containing hundreds of previously published findings.In fact, this functionality has been around for several years in the domain ofapplied psychology through the metaBUS platform (Bosco, Aguinis, Singh, Field,& Pierce, 2015; Bosco, Field, Larsen, Chang, & Uggerslev, 2020).

The metaBUS approach is, essentially,the opposite of the current search process. That is, the platformsystematically extracts all findings from papers *and then* allows users to specify phenomena of interest to retrieveand summarize findings. One clear benefit of the metaBUS approach isportability: The effort involved in extraction can be leveraged to facilitateresearch in perpetuity, even long after achieving your presumed end goal, suchas publication.

With more than 1,100,000 findings andgrowing, metaBUS is, to our knowledge, the largest manually curated collectionof research findings across the social sciences.

The Backstory

As a doctoralstudent, I worked with faculty who were enthusiastic about content-analysis andmeta-analysis in the organizational sciences (e.g., Aguinis, Dalton, Bosco,Pierce, & Dalton, 2011; Dalton, Aguinis, Dalton, Bosco, & Pierce,2012). I began to convince myself that if allprimary studies' research findings were in one big database,conducting and updating meta-analyses would be a breeze. Furthermore, many"science of science" questions could be addressed. Why hadn't we developed sucha database? Probably because the task was perceived as overwhelming. In fact, Iwas repeatedly advised by colleagues to "earn tenure before launching a crazyproject like that!"

But I couldn't help myself. I was ahuge Microsoft Excel enthusiast (still am). And, luckily, I worked in a fieldin which effect sizes were, in more than 90% of cases, correlations reportedwithin nice, neat matrices. There's a lot of data in those matrices—zero-ordereffects ripe for meta-analyzing. In 2011, to take my mind off my dissertation,I started experimenting with combinations of PDF extraction software and VisualBasic for Applications (VBA) script in Excel (the script was needed totranspose the extracted data, remove non-numeric characters, and the like).Eventually, I had the semiautomated extraction process down to roughly 15 to 30seconds per matrix.

The next hurdle was the "vocabularyproblem"; that is, many terms refer to the same thing (Furnas, Landauer, Gomez,& Dumais, 1987). After collecting a few thousand rows of data, it becameclear that the variety of terms was debilitating. I couldn't just search for"performance" (a common applied-psychology topic) and expect to locate resultscomprehensively. I was going to need a hierarchical taxonomy containingessentially all *things* studied in thescientific space, with each entry tagged to the taxonomy. Between 2011 and2013, I collaborated with Kulraj Singh and James Field to develop a taxonomycontaining roughly 5,000 variables/constructs. We also refined protocols forsemiautomated correlation-matrix extraction and manual coding. Together, by2013, we had amassed a database of roughly 200,000 findings and had startedanswering "big science" questions with it (see Bosco et al., 2015). All theoriginal components remain in use by metaBUS today.

Over time, with funding assistanceand new team members, the database grew and evolved. We also experimented witha variety of web-based software platforms before ultimately deciding on aplatform built in R Shiny (Chang, Cheng, Allaire, Xie, & McPherson, 2019).Today, the metaBUS platform is used regularly in the field of appliedpsychology to conduct and assist meta-analyses.

metaBUS: Under the Hood

The process of extracting and classifying findings with metaBUS begins with the semiautomated extraction of correlation matrix content using optical-character-recognition software. The extraction is considered *semi*automated because trained coders oversee the process for each matrix to ensure accurate transcription into a standardized database format.
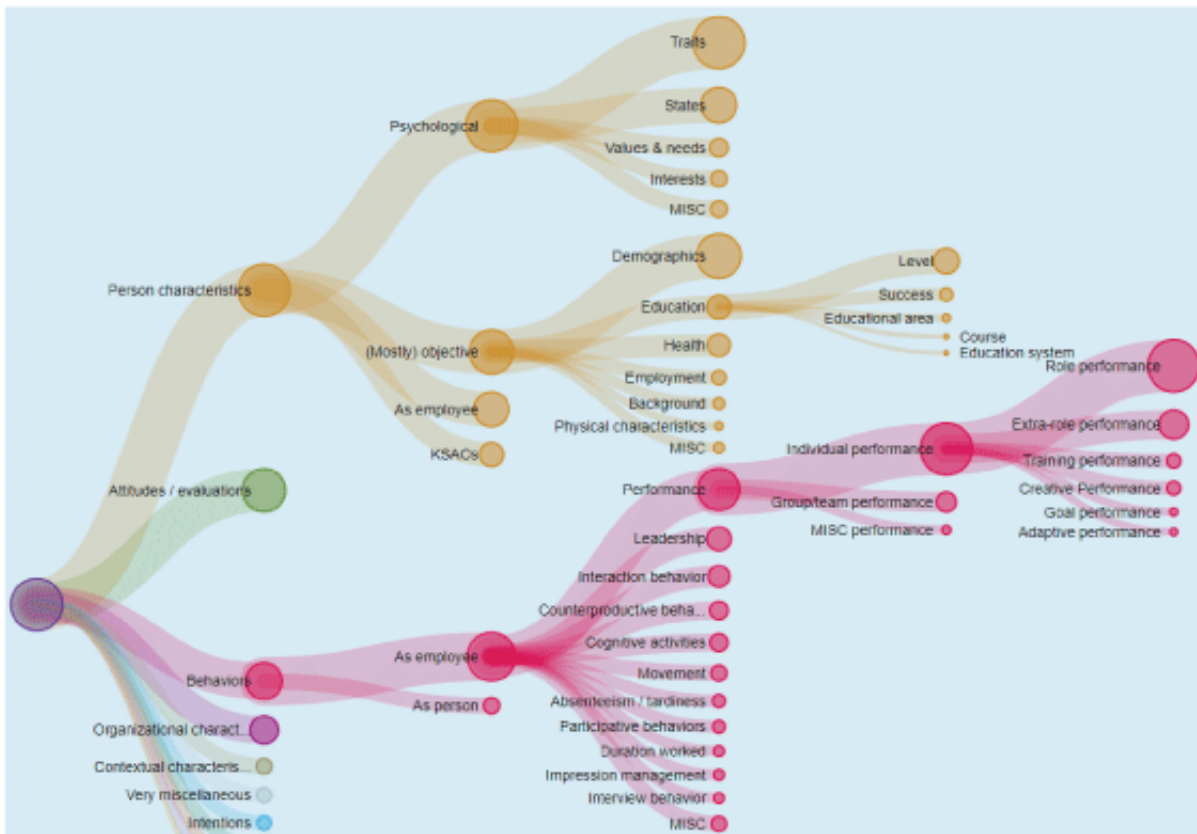
Abbreviated view of the metaBUS taxonomy. The complete taxonomy contains approximately 5,000 nodes and an interactive version is available for viewing at shiny.metaBUS.org.

Next, for each variable contained ineach matrix, trained coders manually assign classifications pertaining to severalattributes (e.g., country of origin, sample size, sample type, response rate).They also manually classify each variable according to the metaBUS taxonomy, ahierarchical "map" of nearly 5,000 constructs and variables studied in appliedpsychology. MetaBUS also leverages R software packages to analyze searchresults, thus enabling instant, rudimentary meta-analyses.

Test-Drive Demonstrations

Following areeasy-to-replicate functionality demonstrations geared toward the first-timeuser—a sort of metaBUS test drive. I demonstrate the metaBUS taxonomy and theprocess of conducting an instant meta-analysis. I also describe how one mightleverage each deliverable to enhance research and graduate training.

Please keep in mind that the metaBUSplatform currently contains findings from appliedpsychology (i.e.,

industrial-organizational psychology and relatedfields such as human-resource management and organizational behavior) simplybecause that's where the project began. There is the potential to expand metaBUSto other areas of psychology, but a search at this time for terms such as"Stroop effect" or "word frequency" or "implicit attitudes toward sugared softdrinks" is likely to be disappointing.

Demonstration 1: Explore the Taxonomy

As a startingpoint for gaining familiarity with metaBUS, I suggest exploring the taxonomy ofconstructs. The taxonomy is arranged hierarchically and contains nearly 5,000construct or variable names that appear in applied psychology. As shown in thefigure on the opposite page, the taxonomy starts with broad branches (e.g., *behaviors*, *attitudes*), each of which subsumes finer level topics. Forinstance, *behaviors* subsumes *employee performance* and *absenteeism.* During data collection,coders manually assign all reported variables a taxonomic node (represented bya five-digit unique identifier), providing some degree of shelter from thevocabulary problem.

Research use cases: Researcherscan explore the taxonomy to browse and ascertain the variety of terms used torefer to a given construct. As an example from applied psychology, thefollowing terms that have appeared in published sources are equivalent: *turnover intention*, *quit intention*, *convictionof decision to quit*, *considering exit*, *perceived chance of leaving*. Thus, fora traditional literature search, such terms could be added to otherto-be-searched variants. Alternatively, to conduct a metaBUS literature searchindependent of text-string shortcomings, one could rely on taxonomy codes inaddition to—or instead of—text strings. For example, rather than dealing withthe variety of terms used to refer to turnover intention, you could simplyspecify *20179* (i.e., an arbitrary5-digit code representing quit intentions in the metaBUS taxonomy), which willcapture all entries tagged by coders as *20179*during data collection, as well as all the node's "children."

Education use cases: Students caninteract with the taxonomy to gain a rapid awareness of the universe ofconstructs studied within an entire scientific space along with an understandingof constructs' conceptual neighbors. Furthermore, because the node size isscaled according to the construct's frequency of occurrence in tables ofresults (larger nodes are studied more frequently), students can quickly gainan understanding of variance in research attention across topics, indicating,for example, areas ripe for meta-analysis.

My colleagues and I urge readers tobrowse the metaBUS taxonomy and then, as a thought experiment, consider ascientific field with which they have little familiarity. In my case, such afield is sociology. I asked myself, *Whatwould an analogous map of sociology look like*? I have no idea, but I canimagine the enormous educational value of a "map" of sociology topics, arrangedby frequency, especially if I were a student of sociology.

Demonstration 2: Conduct an Instant Meta-Analysis

When presentingmetaBUS at conferences, I often demonstrate the instant meta-analysis featureby asking the audience to name two constructs. (It feels a bit like performinga magic trick: "Pick a construct, any construct.") On one occasion, theaudience picked *turnover intention* and*self-efficacy*. I submitted the twoconstructs to a metaBUS search and, abracadabra, within a few seconds, Ipresented the results of an instant, rudimentary meta-analysis that returned 49effects with a mean *r* of –.06. Trythe same analysis

by duplicating the search terms shown in the figure on thispage on the metaBUS platform (shiny.metaBUS.org).

To conduct an instant meta-analysis, two search terms must be specified in the form of taxonomic codes, text strings, or a combination of the two. (See Bosco et al., 2020 for new "exploratory" meta-analysis search functionality that requires specifying only one search term.) Because the taxonomy contains roughly 5,000 elements, the number of possible taxonomic search-term permutations is quite large. However, many possible pairs return zero results; thus, we urge users to try the platform at first by considering searches that involve popular research topics. (As noted earlier, construct frequency may be ascertained by contrasting the metaBUS taxonomy's node sizes.)

Screenshot of the metaBUS instant meta-analysis feature. Relation between turnover intention and self-efficacy displayed. Cloud-based version available at shiny.metabus.org/q/da73b175-0b8e-4349-b4a5-9076edd09690.

Research use cases: The most obvious use case of the metaBUS instant meta-analysis feature is to facilitate the location of research findings. Indeed, findings are often difficult to locate because their variables may have played an ancillary role in the study (e.g., age as a control variable) and are not visible in the article's abstract or keywords. Thus, researchers conducting meta-analyses may wish to search metaBUS to locate findings that may have gone overlooked following a traditional literature search. Furthermore, authors frequently require examples of previous findings to craft manuscript introductions or to justify hypothesis statements. In situations like these, a search engine of findings such as metaBUS becomes highly useful.

Additionally, we envision a varietyof meta-scientific studies using the metaBUS database. For example, studiesthat have examined the frequency distribution of p values (i.e., p-curves)have lamented existing reporting conventions (e.g., $p < .05$) that prevent the calculation of exact p values. In contrast, the metaBUSdatabase contains more than 1,000,000 findings, each with a correspondingsample size. This would allow the largest investigation on the p-curve, with available nuance bypublication year and

bivariate relation type, and the analyses could likely becompleted in a single day! Again, the up-front effort associated with theextraction of findings is portable to other purposes.

Education use cases: Meta-analyses have increased in popularity over the last several decades. In fact, many graduate degree programs now include meta-analysis training. The metaBUS instant meta-analysis feature could be an ideal teaching tool for these courses. Indeed, using the platform, users can interact with the contributing effects (i.e., toggle include/exclude and view impact on summaries), limit by sample size, publication year, and the like. Furthermore, students could gain familiarity with scripting languages by viewing the R script used to derive the meta-analytic estimate.

A Well-Funded, Collaborative Future

At this point, onemight wonder, "can metaBUS be adapted to psychology more broadly—beyond appliedpsychology?" My knee-jerk response is, "yes, with sufficient elbow grease andcollaboration."

As my colleagues and I describe in *AMPPS* (2020), the contents ofcorrelation matrices are easy to extract and contain a wealth of information.However, effect-size reporting in other areas (e.g., much of experimentalpsychology) is precisely the opposite—the articles often contain relatively feweffects, and the sheer variety of effect-size indices (and conditions that giverise to them) makes challenging the development and navigation of codingprotocols. Put differently, efficient coding of experimental research ispossible, but it will require some serious head-scratching. However, there arecertainly other areas of psychology (e.g., military, personality,developmental, educational) that have traditionally reported correlationmatrices or, over time, have done so with increased frequency. In my view, thebest chance psychology has to demonstrate and promote effective, accuratecuration of the research backlog in the near term, is to (a) focus on areasthat have reported correlations by convention (i.e., worry about experimentalresearch later on), (b) secure long-term funding for *manual* curation, and (c) assume a collaborative mindset.

I will elaborate.

Regarding funding, I love a magicpill as much as the next person. However, I believe it is unwise, at this time,to sink millions of dollars into unattended algorithm-based/artificialintelligence approaches with hopes of achieving a one-click solution tocurating psychological research. In my view, it would be far wiser to fundmanual effort-driven curation at first, for these reasons:

It would cost less to pay graduatestudents (e.g., during the summer) to manually code—or perhaps double-code—thebacklog.

This effort would result in a morereliable database (compared with automated approaches), and its contents wouldbe immediately useful for a variety of scientific purposes.

The data created from manual effortswould serve as the ideal "answer key" for *later*development of automated solutions.

Viewed through this lens, it wouldseem funders and researchers often put the cart before the horse and wasteresources. So, for the near term, and to put it frankly, perhaps we shouldleave the engineering work

to engineers and the psychology work topsychologists—at least until our data warehouse is in order. Regardless,dedicated, long-term funding should flow directly from psychology's knownstakeholders to develop and maintain these manually curated search engines. Toprioritize the curation and indexing of one's corpus of findings is, in myopinion, a no-brainer.

Regarding collaborative mindsets, letme first note that competition is not necessarily an unhealthy thing. Indeed,academics regularly compete for limited journal space and research funding; weare not strangers to the recognition economy and relish being the "first," the"founder," the "lead," the "principal investigator," and the like. However,some efforts, such as those involved in curating an entire scientific field'sfindings, rely on coordinated efforts (e.g., Collins, Morgan, & Patrinos,2003). Fewer resources would be wasted by coordinated effort than a dozenego-driven teams hoping to secure their respective places as top gun. Thiswould only lead to delayed progress or, worse, incompatible database formatsthat limit the ability to merge data and answer larger questions.

I am optimistic that the metaBUS approach (by any other name) will spread throughout psychology and beyond. Indeed, research findings are important, and it's unwise to bet against the value provided by search engines. I just hope that it will be done collaboratively so that efficiency and scientific insight might prevail, and so that scientists—for the near-term, at least—resist the urge to develop "magic pills." This is a job that, at the present time, requires elbow grease for proper completion.

metaBUS at a Glance

• MetaBUS is a cloud-based research-synthesis platform sittingatop the world's largest collection of curated social-science research

findings.

• More than 1,113,000 findings from more than 14,000 publishedarticles

• More than 1,175 users (faculty, students, and practitioners) whohave together conducted roughly 13,000 searches

• Funding sources: None until 2014; later sources include theNational Science Foundation and the Social Sciences Humanities Research Council

• 44 team members (listed at metabus.org/about_us/)

• Learn more at metabus.org, or explore the platform,including an interactive version of the taxonomy, at shiny.metabus.org. For a more detailed look at metaBUS, also see Boscoet al. (2020).

References

Bosco, F. A., Aguinis, H., Singh,K., Field, J. G., & Pierce, C. A. (2015). Correlational effect sizebenchmarks. *Journal of Applied Psychology*,*100*, 431–449.

Bosco, F. A., Field, J. G., Larsen,K., Chang, Y., & Uggerslev, K. L. (2020). Advancing meta-analysis

withknowledge management platforms: Using metaBUS in psychology. *Advances in Methods and Practices inPsychological Science*. Advance online publication.doi:10.1177/2515245919882693

Chang, W., Cheng, J., Allaire, J.,Xie, Y., & McPherson, J. (2019). shiny: Web application framework for R(Version 1.3.2) [Software]. Retrieved fromhttp://cran.r-project.org/web/packages/shiny

Cohen, J. (1988). *Statistical power analysis for thebehavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.

Collins, F. S., Morgan, M., &Patrinos, A. (2003). The Human Genome Project: Lessons from large-scalebiology. *Science*, *300*(5617), 286–290.

Furnas, G. W., Landauer, T. K.,Gomez, L. M., & Dumais, S. T. (1987). The vocabulary problem inhuman-system communication. *Communicationsof the ACM*, *30*, 964–971.

Neppali, K., Caragea, C., Mayes,R., Nimon, K., & Oswald, F. (2016). MetaSeer.STEM: Towards automatingmeta-analyses. In D. Schuurmans &M. Wellman (Eds.), *Proceedings of theThirtieth AAAI Conference on Artificial Intelligence and the Twenty-EighthInnovative Applications of Artificial Intelligence Conference* (Vol. 5,pp. 4035-4040). Palo Alto, CA:Association for the Advancement of Artificial Intelligence.