INTERSPEECH 2017: Show & Tell Contribution
August 20–24, 2017, Stockholm, Sweden

MetaLab: A repository for meta-analyses on language development, and more

Sho Tsuji[1], *Christina Bergmann*[2,3], *Molly Lewis*[4], *Mika Braginsky*[5], *Page Piccinini*[6,3], *Michael C. Frank*[7], *Alejandrina Cristia*[2,3]

[1]University of Pennsylvania, USA
[2]Laboratoire de Sciences Cognitives et Psycholinguistique (ENS, EHESS, CNRS)
[3]Département d'Etudes Cognitives, Ecole Normale Supérieure, PSL Research University, France
[4]Computation Institute, University of Chicago, USA
[5]Department of Brain and Cognitive Sciences, MIT, USA
[6]Laboratoire de NeuroPsychologie Interventionnelle (ENS, EHESS, CNRS)
[7]Department of Psychology, Stanford University, USA

tsujish@gmail.com, chbergma@gmail.com, mollylewis@uchicago.edu, mika.br@gmail.com, page.piccinini@gmail.com, mcfrank@stanford.edu, alejandrina.cristia@ens.fr

Abstract

MetaLab is a growing database of meta-analyses, shared in a github repository and via an interactive website. This website contains interactive tools for community-augmented meta-analyses, power analyses, and experimental planning. It currently contains a dozen meta-analyses spanning a number of phenomena in early language acquisition research, including infants' vowel discrimination, acoustic wordform segmentation, and distributional learning in the laboratory. During the Show and Tell, we will demonstrate how to use the online visualization tools, download data, and re-use our analysis scripts for other research purposes. We expect MetaLab data to be particularly useful to researchers interested in early speech perception. Additionally, the infrastructure and tools can be adopted by speech scientists seeking to perform and utilize (meta-)meta-analyses in other fields.

Index Terms: language acquisition, infancy, childhood, speech perception, meta-analyses, visualization, database, experiment planning, big data

1. Introduction

Recent years have seen a surge of concern for the replicability and generalizability of results from psychological experiments, including those of the type often used to study speech perception [1]. One way to address such issues involves the use of meta-analysis [2], a statistical tool to aggregate over a whole body of empirical evidence that is useful in at least three ways. First, meta-analyses can be used to estimate the effect size associated with a phenomenon collapsing across variation in methodology and conceptualization thought to be irrelevant. Second, meta-analyses allow a principled measure of the effects of conceptual and methodological factors on effect size. Third, by virtue of compiling evidence from all studies on a given phenomenon, meta-analytic data can also be interrogated to "diagnose" a field for inappropriate research and publication practices. For instance, meta-analysts can check whether studies are appropriately powered and whether there is evidence that publication is biased against non-significant findings and/or certain directions of results.

For all of these reasons, meta-analyses are extremely useful for members of a community pursuing empirical research. However, they remain marginal due to three main roadblocks. First, they are challenging to build, as considerable effort is required to compile information. Second, they remain of limited use when they appear solely as a paper publication, for instance in the form of a table that cannot be easily re-analyzed, and which in any case "ages" with the published paper. Third, the statistical expertise required to profit from meta-analytic data, although not extensive, does require some personal investment. In previous work, we have proposed the concept of community-augmented meta-analyses to divide the effort of creating and updating meta-analyses [3], and which ameliorates the first two roadblocks. With MetaLab, we attack the last roadblock, bringing the power of meta-analyses within more general reach.

2. The concept and tools

We aim to enable the research community to more easily use and re-use meta-analyses. There are literally hundreds of textbooks, and probably many online courses introducing the general method. There are also a number of general guidelines (e.g., [4]). However, to our knowledge, there is no resource that combines a hands-on introduction to meta-analyses, with easy access to the basic tools allowing a wide range of users to benefit from them.

MetaLab is currently implemented as a point-and-click shiny webapp `metalab.stanford.edu`, containing a number of analysis and visualization tools, all of them implemented within the open-source, multi-platform statistical computing software R [5]. To begin with, tools include a range of analyses involved in meta-analysis proper, mainly following [6]. These include a range of meta-analytic models (e.g., using multilevel grouping studies reported in the same paper as well as empirical Bayes estimations), and standard visualizations (violin plots of the effect size distribution, forest and funnel plots). For these, we rely greatly on the metafor package [7].

We have also gone beyond these basic tools bearing the empirical researcher in mind. Specifically, we include experimental planning tools, which use the meta-analytic effect size retrieved from a given meta-analysis to inform prospective power analyses. This allows researchers interested in a similar topic to make more informed decisions regarding the number of infants they should plan on testing, avoiding later disappointments common when data are sparse. An "experiment simulation" tool further allows web users to play around with key param-

Table 1: *Sample of the meta-analyses currently available from MetaLab.*

Source	Description	Papers (expts)
Cristia (in rev.)	Learning of phonetic and phonotactic generalizations from a short exposure in the laboratory	24 (57)
Tsuji & Cristia (2014)	Discrimination of native and non-native vowels	44 (162)
Bergmann & Cristia (2015)	Recognition of briefly familiarized words from running, natural speech.	68 (285)
Lewis & Frank (in prep.)	Mutual exclusivity (i.e., look at novel object when hearing a novel word)	20 (60)
Lammertink et al. (2016)	Sound symbolism (e.g., look at a round object when hearing words like "bouba")	11 (44)
Lewis & Long (unpublished)	Label advantage in visual categorization	14 (49)
Frank, Lewis, & MacDonald (2016)	Recognition of familiar words	6 (14)

eters and observe the potential outcomes of their study.

Finally, our code also includes tools to generate a number of reports bearing on questions of theoretical and empirical interest that will interest those working on language acquisition. For example, on the theory side, we have investigated whether linear versus higher-order models best accommodate changes found with age across each and all meta-analyses. On the empirical side, we have implemented p-curve analyses [8].

All of these analysis and visualization tools can be used through an intuitive point-and-click interface offered by the online platform. Additionally, those interested in digging deeper into our code and data, or those interested in generating a similar resource centered on other topics, would benefit from a visit to our github repository.

3. Current datasets and standards

Currently, MetaLab contains 887 effect sizes across meta-analyses in a range of phenomena within language acquisition, based on data from 252 papers collecting 11363 infants aged between 0 and 36 months (see Table 1 for a sample). While several of these meta-analyses pre-dated the creation of MetaLab, others were developed at the same time as this resource, and all of them have been modified and improved in the course of this collaborative experience. Specifically, building this common resource with a more general outlook allowed us to identify which dimensions any meta-analysis bearing on infant perception must absolutely contain, and which would only be interesting to researchers working on a particular phenomenon. As a result, we were able to define the structure of any meta-analysis that could be potentially included (available as a yaml file). Most saliently, we have learned to pay attention to the complexities of data sets, which sometimes contain between-participant comparisons, others within-participant comparisons, and yet others repeated wihin-participant comparison contributing to different levels of analysis.

As with the tools, there are options for "advanced" users to get even more from MetaLab meta-analyses. In particular, the original data for each meta-analysis can be downloaded from MetaLab, thus allowing users to explore additional dimensions that have been coded but not directly included in MetaLab. Additionally, we think that others who may be interested in building a MetaLab-like resource in a different domain could draw inspiration from our data structure.

4. Future directions

Thanks to funding from SSMART, MetaLab continues growing, as we update and add meta-analyses. We are also keenly interested in promoting the use of MetaLab and meta-analyses more generally. We therefore have already developed a set of documents geared at potential contributors, as well as a set of videos introducing MetaLab and meta-analyses on youtube.

In several manuscripts, we are exploring the possibility of combining information across several meta-analyses to answer theoretical questions spanning across phenomena - i.e., to constrain theorizing. To this end, we are assessing when and how one can compare effect sizes across meta-analyses, taking meta-analytic effects as indices of underlying skills. For instance, we believe that, at least based on currently available evidence, phonological categories seem to develop prior to, and potentially independently from, the skills involved in recognizing words from running speech [9], and more generally that current experimental data support parallel or interactive, but not stage-like, developmental theories [10].

5. Acknowledgements

The present work was supported by the Berkeley Initiative for Transparency in Social Sciences and the Laura and John Arnold Foundation; the H2020 European Research Council [Marie Skłodowska-Curie grant Nos 660911 and 659553]; and the Agence Nationale pour la Recherche [ANR-14-CE30-0003 MechELex, ANR-10-IDEX-0001-02 PSL*, ANR-10-LABX-0087 IEC].

6. References

[1] O. S. Collaboration *et al.*, "Estimating the reproducibility of psychological science," *Science*, vol. 349, no. 6251, p. aac4716, 2015.

[2] R. Rosenthal and M. R. DiMatteo, "Meta-analysis: Recent developments in quantitative methods for literature reviews," *Annual review of psychology*, vol. 52, no. 1, pp. 59–82, 2001.

[3] S. Tsuji, C. Bergmann, and A. Cristia, "Community-augmented meta-analyses: Toward cumulative data assessment," *Perspectives on Psychological Science*, vol. 9, no. 6, pp. 661–665, 2014.

[4] D. Moher, A. Liberati, J. Tetzlaff, D. G. Altman, P. Group *et al.*, "Preferred reporting items for systematic reviews and meta-analyses: the prisma statement," *PLoS med*, vol. 6, no. 7, p. e1000097, 2009.

[5] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2015. [Online]. Available: http://www.R-project.org

[6] L. Hedges and I. Olkin, *Statistical analysis for meta-analysis*. Academic Press, 1985.

[7] W. Viechtbauer *et al.*, "Conducting meta-analyses in r with the metafor package," *J Stat Softw*, vol. 36, no. 3, pp. 1–48, 2010.

[8] U. Simonsohn, L. D. Nelson, and J. P. Simmons, "P-curve: a key to the file-drawer." *Journal of Experimental Psychology: General*, vol. 143, no. 2, p. 534, 2014.

[9] C. Bergmann, S. Tsuji, and A. Cristia, "Top-down versus bottom-up theories of phonological acquisition: A big data approach," *OSF Preprints*, 2017. [Online]. Available: https://osf.io/vypwu

[10] M. Lewis, M. Braginsky, S. Tsuji, C. Bergmann, P. E. Piccinini, A. Cristia, and M. C. Frank, "A quantitative synthesis of early language acquisition using meta-analysis," Jan 2017, preprint available on PsyArXiv. [Online]. Available: osf.io/pcqac/