Using automation to combat the replication crisis: A case study from controlled-rearing studies of newborn chicks

Samantha M. W. Wood¹✉ and Justin N. Wood¹

¹University of Southern California, Department of Psychology, Building SGM Room 501, 3620 South McClintock Ave., Los Angeles, CA 90089

The accuracy of science depends on the precision of its methods. When fields produce precise measurements, the scientific method can generate remarkable gains in knowledge. When fields produce noisy measurements, however, the scientific method is not guaranteed to work—in fact, noisy measurements are now regarded as a leading cause of the replication crisis in psychology. Scientists should therefore strive to improve the precision of their methods, especially in fields with noisy measurements. Here, we show that automation can reduce measurement error by 60% in one domain of developmental psychology: controlled-rearing studies of newborn chicks. Automated studies produce measurements that are 3-4 times more precise than non-automated studies and produce effect sizes that are 3-4 times larger than non-automated studies. Automation also eliminates experimenter bias and allows replications to be performed quickly and easily. We suggest that automation can be a powerful tool for improving measurement precision, producing high powered experiments, and combating the replication crisis.

Correspondence: samantha.m.w.wood@gmail.com

The scientific method has produced astonishing gains in human knowledge. In the past decade, however, it has become clear that the scientific method does not always work, producing large bodies of unreproducible findings (1). In social psychology, for example, rigorous replication studies have failed to replicate many published findings, including textbook phenomena such as ego depletion (2) and social priming (3). Likewise, in a large-scale replication study of 100 experiments from three high-ranking psychology journals, only one-third to one-half of the original findings were observed in the replication studies (4). Thus, some findings in psychology are reproducible, whereas others are not. What distinguishes reproducible from non-reproducible findings, and how can we improve our methods to produce more reproducible results?

The field of metascience—the scientific study of science itself—has identified a number of factors contributing to the replication crisis. The most prominent factors are underpowered studies (low sample size relative to effect size), high measurement error (noisy measurements), small effect sizes, experimenter bias, analytic flexibility (also known “P-hacking” and “the garden of forking paths”), and conflicts of interest (1, 5–7). Producing accurate and reproducible research therefore requires combating these factors.

Here, we present a case study of how automation can be used to address many of these factors contributing to the replication crisis. Specifically, we focus on controlled-rearing studies of newborn chicks—an area with deep historical roots in developmental psychology. Our paper is organized as follows. First, we briefly describe controlled-rearing studies and their importance for developmental psychology. Second, we sampled a range of non-automated controlled-rearing studies with newborn chicks and found that these studies tended to produce noisy measurements and small effect sizes—two leading factors in the replication crisis. Third, we show how automation can solve this problem, by allowing newborn chicks to be observed and tested for long periods of time. We found that increasing the amount of data collected from each chick (via automated tracking software) reduced measurement error by 60% and increased the effect size of experiments by about 300%. Automation also eliminates experimenter bias and allows replications to be performed quickly and easily. We conclude that automation can significantly improve the power and accuracy of controlled-rearing studies with newborn animals.

1. Using controlled rearing to explore the origins of intelligence

Controlled-rearing studies are designed to explore the role of experience in shaping perception, cognition, and behavior. Despite widespread interest in this topic, there is little agreement on fundamental issues, including which mechanisms exist in newborn brains and how those mechanisms transform sensory input into knowledge. On one view, newborn animals are equipped with a collection of domain-specific mechanisms, each shaped by evolution to perform a particular function. Some researchers argue that newborns have numerous domain-specific systems (e.g., 8, 9), while others argue that newborns have a smaller number of specialized core knowledge systems (i.e., systems for representing objects, actions, number, and space; 10, 11). According to this class of “nativist” theories, newborns have domain-specific learning mechanisms at the onset of post-natal experience.

On the second view, newborn brains contain flexible and adaptive domain-general learning mechanisms for building domain-specific knowledge through experience (12–14). This view originally traces back to Enlightenment thinkers such as Locke (15) and Hume (16). More recent accounts from computational neuroscience (e.g., 17–19) argue that

the representational abilities of the brain are built from the dynamic interaction between neural growth mechanisms and environmentally produced neural activity. This class of “empiricist” theories predict that domain-specific systems emerge from domain-general learning mechanisms as newborns acquire experience with the world.

Modern psychologists generally accept that newborn brains are equipped with both domain-specific and domain-general learning mechanisms. Nonetheless, there is still little consensus on fundamental issues: researchers disagree about what these mechanisms are, what they consist of, what causes them to emerge, and how they change with experience. How can we advance this debate and characterize the learning mechanisms in newborn brains? As scientists have long noted, human infants are not an optimal population for addressing these questions. When infants participate in experiments, they have already been shaped by days, weeks, or months of experience with the natural visual world. Consequently, for almost any ability found in young human infants, it is not possible to determine whether that ability is hardwired by natural selection or learned from experience. Indeed, there is growing evidence from neuroscience that natural sensory experience plays an important role in building domain-specific knowledge (e.g., for faces and symbols 20, 21). Studies of monkeys also show that brains change rapidly in response to statistical redundancies in the environment, with significant neuronal rewiring occurring in as little as one hour (22, 23). These findings allow for the possibility that even early emerging domain-specific abilities are learned from post-natal experience.

In contrast to studies of human infants, controlled-rearing studies of nonhuman animals can reveal the precise role of post-natal experience on development (24–26). By systematically manipulating the experiences provided to newborn subjects and observing the effects of those manipulations on behavior, we can distinguish the experiences that are causally related to developmental change from those that are not. Controlled-rearing experiments therefore provide an experimental avenue for probing the role of experience in the development of mental abilities.

For example, recent automated controlled-rearing studies have started to reveal the role of visual experience in building object concepts. These studies indicate that newborn animals (domestic chicks) can begin building abstract (view-invariant) object concepts at the onset of vision (27, 28). From an engineering perspective, this is an impressive computational feat: computer vision systems typically require thousands to millions of diverse training images to build abstract object concepts. However, this ability only emerges when newborn chicks are raised with objects that move slowly and smoothly over time (29, 30). When objects move too quickly or non-smoothly, newborn chicks build inaccurate object concepts. Thus, to learn how to perceive objects, chicks need experience with objects that adhere to the spatiotemporal properties of objects in the real world, akin to human children (31). From a mechanistic perspective, these results indicate that newborn brains contain unsupervised tem-

poral learning mechanisms, which leverage the spatiotemporal characteristics of natural visual environments to build object concepts. This machinery has been sculpted by evolution to generate abstract object concepts, but only when the brain receives input from a natural visual world.

Newborn chicks as a model system for studying perceptual and cognitive development. Newborn chicks (*Gallus gallus*) are uniquely suited for studying the development of perception and cognition. Unlike commonly-used animal models in psychology (e.g., rats, pigeons, and monkeys), chickens are a precocial species (mobile in the first day of life) and can be raised in strictly controlled environments immediately after hatching. Newborn chicks can be raised for weeks in environments containing no real-world objects or agents (27). Thus, studies of newborn chicks allow us to test critical hypotheses in the nativist-empiricist debate. With chicks, we can discover mechanisms that are present in newborn brains, and explore how those mechanisms are shaped by experience. While mechanisms found in chicks may or may not have also evolved in humans (the last common ancestor of humans and chickens lived around 310 million years ago), these studies provide an existence proof of the types of mechanisms that are present in newborn brains.

Moreover, recent work from avian neuroscience indicates that cortical mechanisms are largely conserved across birds and mammals (reviewed by 32, 33). At the circuit-level, avian and mammalian brains contain homologous cortical circuits for processing sensory information; these circuits share similarities in terms of cell morphology, the connectivity pattern of the input and output neurons, gene expression, and function. At the macro-level, avian and mammalian brains share the same large-scale organizational principles: both are modular, small-world networks with a connective core of hub nodes that includes prefrontal-like and hippocampal structures. Since cortical mechanisms are shared across birds and mammals, studies of chicks can reveal general characteristics of vertebrate brain development. While chicks have smaller brains than humans, this can provide an advantage to researchers. When attempting to understand a particular phenomenon, it is often valuable to use the simplest system that demonstrates the properties of interest. The fields of neuroscience and biology rely heavily on this strategy—for example, researchers use sea slugs to study the physiological basis of memory storage in neurons (34), worms to study the mechanisms of molecular and developmental biology (35), and fruit flies to study the mechanisms of genetics (36). In a similar vein, controlled-rearing studies of newborn chicks can provide an important (and unique) perspective on perceptual and cognitive development. Next, we briefly describe non-automated controlled-rearing studies with newborn chicks, with a focus on studies that had a significant influence on theories of cognitive development.

2. Non-automated studies of newborn chicks

Controlled-rearing studies of newborn chicks have deep historical roots in developmental psychology (e.g., 37, 38). The

Fig. 1. Scatterplots and boxplots of the (A) measurement error and (B) effect sizes from samples of automated (blue points) and non-automated (red points) controlled-rearing studies. Each point represents the (A) standard deviation or (B) Cohen's *d* from a single condition (only statistically significant conditions are included). The boxplots show the range from 25th to 50th percentile and from 50th to 75th percentile. Points with an asterisk are beyond the range of the graph (greater than 50% standard deviation or greater than 2.5 Cohen's *d*). Across a wide range of studies, the effect sizes obtained with automated methods were much larger than the effect sizes obtained with non-automated methods. Automated methods also produced far more precise measurements than non-automated methods. (C) The number of subjects needed to achieve 80% power for a range of true population performance values for studies with low measurement error (standard deviation = 10%) and high measurement error (standard deviation = 33%). These standard deviations match those from our samples of automated and non-automated studies, respectively. Low measurement error massively reduces the number of subjects needed to achieve adequate experimental power. When measurement error is high, small decreases in true population performance require large increases in the number of subjects needed to achieve 80% power.

majority of these studies were designed to explore whether chicks have domain-specific knowledge in the absence of post-natal experience with the domain.¹ In support of this premise, researchers have reported that newborn chicks have early emerging preferences for faces (39), biological motion (40), and self-propelled motion (41), all without prior experience with faces, biological motion, and self-propelled motion. These findings have been interpreted as evidence for an innate (non-learned) mechanism for detecting animate agents. There have also been reports that newborn chicks have early emerging capacities for object cognition (42) and numerical cognition (43), and that newborn ducklings can perform abstract relational reasoning (44). Together, these studies have been interpreted as strong evidence for nativist theories of cognitive development (10, 11, 38). Since there was little to no opportunity for the animals in these studies to learn about objects, agents, and number prior to testing, empiricist theories of cognitive development cannot readily explain these findings. Given the importance of these findings, it seems vital to assess the precision of the measurements obtained in the studies. In all of the studies cited above, the chicks were tested in a two-alternative forced-choice task,

within a single session lasting a short period of time (generally 6 minutes). This is a short measurement period, so these studies tended to produce noisy measurements. In a sample of 10 non-automated controlled-rearing studies (Table 1), we found that the average standard deviation (measurement error) was 33%, ranging from 17% to 65% across studies (Figure 1A). To illustrate how noisy of a measurement this is, Figure 2 shows visualizations of distributions from noisy measurements (SD = 33%) versus precise measurements (SD = 10%). Note that for the noisy measurements, a large proportion of the data falls both above and below chance levels, hindering assessment of true population performance.

Furthermore, researchers often conclude that chicks succeeded at a task irrespective of whether the chicks performed above or below chance levels (i.e., both familiarity and novelty effects are considered evidence for an ability), providing increased flexibility to find an effect. Unfortunately, recent work from metascience shows that this combination of factors (noisy measurements and flexibility in the direction of the effect) can easily lead to the perception of false effects (5–7).

When data are noisy, studies often produce estimates of performance that are considerably higher (or lower) than the true population performance (5). These inflated estimates of performance can significantly increase false positives rates, especially when researchers have flexibility in terms of the number of subjects that are tested, the analyses that are performed, and the results that are reported (7). Inflated es-

¹These studies generally do not assume a special ecological niche that would differentiate domain-specific knowledge in chicks from domain-specific knowledge in other species. However, it is worth noting that chickens are a precocial species and may thus differ in some ways from humans. For example, chicks can begin exploring their environment at the onset of post-natal experience. This active motor exploration might play an important role in the development of perceptual and cognitive abilities.

Fig. 2. Visualization of normal distributions of data with high measurement error (standard deviation = 33%) and low measurement error (standard deviation = 10%). These standard deviations match those from our samples of automated and non-automated studies, respectively. Despite having the same mean performance (70%), distributions with low measurement error produce much larger effect sizes than distributions with high measurement error (Cohen's $d = 2.0$ versus $d = 0.6$, respectively). Note that while the low measurement error data has a very narrow distribution, much of the high measurement error data falls both above and below chance level (despite the high average performance).

timates of effect size lead researchers to underestimate the number of subjects needed for appropriate statistical power, resulting in underpowered studies.²

A critic might argue that if studies produce statistically significant results despite having noisy measurements, then the results must be particularly robust (because it was possible to detect the effects despite the noise in the signal). However, researchers typically have so many “researcher degrees of freedom” (flexibility in how they prepare, analyze, and report data) that statistical significance can easily be found even in the absence of underlying effects (7), even without multiple hypothesis testing (45). This problem is further exacerbated when data are noisy (5). In other words, statistical significance conveys very little information when measurements are noisy and researchers have high analytic flexibility (for a detailed discussion see 46).

In sum, non-automated controlled-rearing experiments with newborn chicks tend to produce noisy measurements and small effect sizes. Thus, these studies show two of the factors contributing to the replication crisis (6). When researchers collect a single short trial from subjects and have high analytic flexibility, then the resulting noisy measurements can masquerade as true effects. A method that could produce precise measurements and large effect sizes would therefore improve the accuracy of controlled-rearing studies. In the next section, we describe how automation can be used to produce precise estimates of performance from newborn chicks, resulting in large effect sizes and more powerful experimental designs.

²Effect sizes measure the magnitude of an effect relative to the noise (e.g., for a one-sample t -test, Cohen's $d = (\text{mean} - \mu)/\text{SD}$). Because the denominator to compute d is standard deviation (SD), methods that produce larger standard deviations will have smaller effect sizes. Critically, effect sizes and power are inextricably linked: smaller effect sizes require larger numbers of subjects to attain the same level of power (see Figure 1C).

Fig. 3. Illustration of an automated controlled-rearing chamber. The chambers contain no objects other than the virtual objects projected on the display walls. Using microcameras embedded in the ceiling, we record chicks' behavior 24/7 via automated image-based tracking software.

3. Using automation to increase the precision and power of controlled-rearing experiments

Historically, newborn chicks' behavior has been measured through direct observation by trained researchers. This approach has limitations: a researcher can typically observe only a single subject at a time, and these observations are susceptible to experimenter bias. Experimenter bias can alter subjects' behavior and the measurement of that behavior. In other fields—including physics, astronomy, and chemistry—automation has been invaluable for minimizing experimenter bias in research procedures, while also removing sources of noise by standardizing data collection procedures.

To apply these benefits of automation to controlled rearing research, our lab developed a fully automated controlled-rearing method (27). The method allows newborn chicks to be raised for weeks within automated controlled-rearing chambers. As illustrated in Figure 3, these chambers provide complete control over all visual object experiences. The chambers contain extended surfaces only, and stimuli are presented to the chicks by projecting virtual objects on two display walls (LCD monitors) situated on opposite sides of the chamber. Thus, the chicks' visual object experiences are limited to the virtual objects projected on the display walls.

All stimuli presentation, data collection, and behavioral tracking are performed by computers, allowing the chicks' behavior to be measured continuously (24/7). The chambers track each chick's location at the rate of 9 samples/sec, using micro-cameras and automated image-based tracking software (EthoVision XT, Noldus Information Technology, Leesburg, VA). Image-based tracking operates over a digital recording of the animal's behavior, which maintains an objective view of events, allows analyses to be repeated, and enables researchers to perform exploratory analyses of variables not originally considered.

This method also produces hundreds of hours of data per chick over their first few weeks of life. Accordingly, we can obtain precise estimates of each chick's performance and determine whether each chick succeeded or failed at the experimental task. Moreover, since the entire data collection process is automated, this method eliminates experimenter bias (e.g., bias that may occur when coding the subject's behavior, presenting stimuli to the subject, or deciding whether to include the subject in the final analysis).

To quantify the benefits of automation, we compared the measurement error and effect sizes from representative samples³ of automated and non-automated controlled-rearing studies. For the automated studies, the average standard deviation was 10% (blue points in Figure 1A), whereas for the non-automated studies, the average standard deviation was 33% (red points in Figure 1A). Thus, automated studies produce measurements that are 3-4 times more precise than non-automated studies. Next, we analyzed whether automated studies produce larger effect sizes (Cohen's d) than non-automated studies (for reference, $d = 0.20$ is considered small, $d = 0.50$ medium, and $d = 0.80$ large). We found that automated studies tend to produce very large effect sizes (average Cohen's $d = 1.43$ in our sample, blue points in Figure 1B). In contrast, non-automated studies produce small to medium effect sizes (average Cohen's $d = 0.37$ in our sample; red points in Figure 1B). Thus, automated studies produce effect sizes that are 3-4 times larger than non-automated studies.

To better understand why automation improves measurement precision, we re-analyzed previously published data from our lab. In the analysis, we tested how the inclusion of more data per chick influenced the observed effect size for the first two studies published from our lab (27, 47). As shown in Figure 4, when only 12 minutes of data were analyzed per chick (the first 6 min of 2 trials = 12 min of data per subject), the average Cohen's d hovered in the medium range, $d = 0.5$. When we increased the data to include a full day of these shortened trials (the first 6 min of 24 trials = 144 min of data per subject), the average Cohen's d rose to $d = 1.1$. When we included all of our available data (20 min trials $\times$ 24 trials per day $\times$ 7 days = 3,360 min of data per subject), the average Cohen's d rose to $d = 1.9$. Thus, increasing the amount of data collected per chick reduced the measurement error by a factor of 2.6 and increased the effect size by a factor of 4.

In addition to improving experimental accuracy, another benefit of reducing measurement error is that fewer subjects are needed to detect an effect. While the standard recommendation in response to the replication crisis is to *increase* sample size, statisticians are now suggesting that an even better focus is to improve the quality of measurements (reviewed by 46). Mathematically, reducing measurement error by a factor of 3 is as good as multiplying the sample size by 9.

To illustrate this effect, we computed the sample size needed to obtain 80% power for a range of "true" population performance magnitudes (Figure 1C; 48). Automated methods that produce precise measurements ($SD = 10\%$) require far smaller numbers of subjects for 80% power than

non-automated methods that produce noisy measurements ($SD = 33\%$). For example, achieving adequate power to detect a true population performance of 58% (chance = 50%) with automated methods requires 15 chicks, while detecting the same effect with non-automated methods requires 136 chicks!

While automated controlled-rearing methods cannot be applied to human subjects, the principle of increasing the precision of measurements need not be unique to studies of newborn chicks. Methods that reduce statistical noise can improve the effect sizes and power of infant studies as well. For example, automated methods using tablets and head-mounted eye tracking may allow researchers to collect more data from each subject, while reducing experimenter bias in the data collection process.

4. Automation allows individual-level statistical analyses

After collecting a single, short trial from each chick, experimenters can compute a group-level measure of average performance. However, this group-level statistic does not tell us whether all of the chicks had the ability of interest or whether just a subset of the chicks had the ability. In fact, studies with noisy measurements typically find that a large proportion of the subjects performed in the *opposite direction* of the reported effect (Figure 2). This is problematic from a theoretical perspective because if an ability is foundational to perception and cognition, then it should be present in most—if not all—newborn animals (provided we have sufficiently precise methods to detect the ability).

Automated testing provides a direct solution to this problem. By collecting hundreds of hours of data per chick, automation allows researchers to obtain precise measurements of each subject's performance and assess whether each chick performed above (or below) chance level (e.g., 27, 49). With massive amounts of data from each chick, these individual-level analyses are often highly powered: many reach 5-sigma levels of statistical significance (the statistical threshold for new discoveries in theoretical physics) for individual subjects.

Another benefit of individual-level analyses is that each subject can serve as a replication of an effect. If we consider an experiment as being on an individual level (rather than group level), then each subject provides an opportunity to replicate an effect. Experiments at the individual level are only possible with enough measurements from each subject to reliably reject or fail to reject the null hypothesis for each individual. Thus, automation permits individual-level internal replications.

Individual-level analyses may also be possible in studies of infants and toddlers. Online methods of testing allow children to participate in studies without having to continually travel back and forth to a physical laboratory. By recruiting infants and toddlers for multiple testing sessions, researchers of human development can also examine variation between individuals and use individual-level analyses as internal replications of an effect.

³For our samples, we chose studies that used a two-alternative forced-choice test, sampled behavior during preselected measurement windows, and did not use a training procedure. For the non-automated studies, we selected studies that have had an influence on theories of developmental psychology (e.g., studies arguing for innate capacities for object, agent, and number cognition). For the automated studies, we selected the first 10 automated studies published from our lab with significant results. For both samples, we included all of the conditions reporting statistically significant effects. In the Supplementary Data, we provide an Excel file that includes all of the data and calculations used to produce the measurement error and effect size estimates across the samples.

Fig. 4. Automation can significantly **(A)** increase the effect size and **(B)** decrease the measurement error of an experiment, by allowing large amounts of data to be collected from each subject. For example, for Wood (2013) and Wood (2014), increasing the amount of data collected from each chick increased the observed effect size by a factor of 4 and reduced the measurement error by a factor of 2.6. Thus, collecting more data per subject significantly increases the precision of data, leading to high powered experiments.

5. Eliminating analytic flexibility through pre-registration and automation

Our paper has focused largely on improving measurement error and effect sizes. However, combating the replication crisis also requires reducing analytic flexibility (“researcher degrees of freedom”) in experimental designs. For example, researchers can (1) engage in flexible data collection (e.g., continuing to collect data until the results reach statistical significance), (2) perform a range of analyses and report only those analyses that reached statistical significance, and (3) continue to test a hypothesis with slight methodological changes until a statistically significant result emerges, then frame the failed attempts as limits to the phenomenon (50). These research practices dramatically increase false-positive rates (7).

The most direct way to combat these problems is to pre-register the design and analyses of experiments. Preregistration requires researchers to describe, in as much detail as possible, their plans for the study (e.g., number of subjects, stimulus materials, procedures, measures, rules for excluding data, plans for data analysis, predictions/hypotheses) and to post those plans in a time-stamped, locked file in an online repository that can be accessed by editors and reviewers (and, ultimately, by readers). By forcing researchers to make these decisions before collecting data, preregistration minimizes researcher degrees of freedom and reduces the probability of obtaining false-positives.

Automation can also be useful for combating experimental and analytic flexibility. To provide an example, our lab recently started automating data analysis procedures. We first

create an R script of the analyses we plan to run, then use that script to analyze the raw data from the animal tracking computer. Thus, we formalize the statistical tests we plan to run *before* looking at the data. For each experiment, we then create a “reproducibility package” containing the R script and raw data. Running the package automatically reproduces the statistical analyses from the raw data. We are also starting to automate replication procedures, by creating “replication packages” containing all of the stimuli, materials, and code needed to replicate our experiments. Running the package automatically recreates all of the stimuli presented to the chicks, in exactly the same order as the original experiment. In the future, we plan to begin making these reproducibility and replication packages openly available to researchers (e.g., by posting them on our lab website), helping contribute to a culture of transparent and reproducible research. Automated data analysis procedures could also be used by laboratories that study infants and children to improve the reproducibility of statistical analyses.

6. Conclusion

A central goal in science is to produce robust and reproducible research. One strategy for improving reproducibility is to reduce measurement error (46). In this paper, we show how automation can reduce measurement error by 60% in controlled-rearing studies of newborn chicks. Specifically, automated controlled-rearing studies produce measurements that are 3-4 times more precise than non-automated studies and produce effect sizes that are 3-4 times larger than non-

automated studies.

Our paper focuses on controlled-rearing experiments, but this strategy—which emphasizes improving measurement precision—could also be useful for other areas of developmental psychology. In response to the replication crisis, developmental psychologists have tended to focus on other strategies (e.g., increasing sample size, preregistration, open data policies, and performing exact replications). We agree that all of these strategies are valuable, but we suggest that researchers should be especially focused on reducing the amount of noise in their data.

How can researchers improve measurement precision in their experiments? One straightforward strategy is to collect more data from each subject. We found a strong positive correlation between the amount of data collected from each subject and the precision of the measurements (Figure 4). Moreover, by collecting more data per subject, researchers can determine the homogeneity or heterogeneity of their effects across infant and toddler populations.

Many researchers in developmental psychology are already moving in this direction. For example, while it is not possible to experimentally test human infants 24/7, new automated internet-based platforms, such as “Lookit”, allow families to participate in behavioral studies via webcam (51, 52). Thus, researchers can begin recruiting larger numbers of infants and collecting repeated measurements. There is also growing use of automation to track infants’ eye movements (53) and language acquisition (54). Roy and colleagues, for example, used automation to collect more than 200,000 hours of audio and video recordings from all of the rooms in a child’s house, providing a detailed case study of the emergence of language.

Robust, replicable findings are vital for formulating and evaluating scientific theories. We conclude that automation can be a powerful tool for combating the replication crisis and improving the precision of scientific measurements in developmental psychology. Precise measurements make it easier to detect true effects, reduce the probability of detecting false effects, and decrease the number of subjects needed to achieve adequate power. Automated methods also protect against experimenter bias, enable individual-level statistical analyses, and allow replications to be performed quickly and easily.

Bibliography

1. J. P. Ioannidis. Why most published research findings are false. *PLoS Med*, 2(8):e124, 2005. ISSN 1549-1676 (Electronic) 1549-1277 (Linking). doi: 10.1371/journal.pmed.0020124. Ioannidis, John P A eng 2005/08/03 09:00 PLoS Med. 2005 Aug;2(8):e124. Epub 2005 Aug 30.
2. Martin S Hagger, Nikos LD Chatzisarantis, Hugo Alberts, Calvin O Anggono, Cedric Batailler, Angela R Birt, Ralf Brand, Mark J Brandt, Gene Brewer, and Sabrina Bruyneel. A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*, 11(4):546–573, 2016. ISSN 1745-6916.
3. Richard A Klein, Kate A Ratliff, Michelangelo Vianello, Reginald B Adams Jr, Štěpán Bahnik, Michael J Bernstein, Konrad Bocian, Mark J Brandt, Beach Brooks, and Claudia Chloe Brumbaugh. Investigating variation in replicability: A “many labs” replication project. *Social psychology*, 45(3):142, 2014. ISSN 2151-2590.
4. Open Science Collaboration. Estimating the reproducibility of psychological science. *Science*, 349(6251), 2015. ISSN 0036-8075.
5. Eric Loken and Andrew Gelman. Measurement error and the replication crisis. *Science*, 355(6325):584–585, 2017. ISSN 0036-8075.
6. Marcus R Munafò, Brian A Nosek, Dorothy VM Bishop, Katherine S Button, Christopher D Chambers, Nathalie Percie du Sert, Uri Simonsohn, Eric-Jan Wagenmakers, Jennifer J

- Ware, and John PA Ioannidis. A manifesto for reproducible science. *Nature Human Behaviour*, 1:0021, 2017. ISSN 2397-3374.
7. Joseph P Simmons, Leif D Nelson, and Uri Simonsohn. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological science*, 22(11):1359–1366, 2011. ISSN 0956-7976.
8. Leda Cosmides and John Tooby. Origins of domain specificity: The evolution of functional organization. *Mapping the mind: Domain specificity in cognition and culture*, pages 85–116, 1994.
9. Steven Pinker. *The blank slate: The modern denial of human nature*. Penguin, 2003. ISBN 0142003344.
10. S. Carey. *The origin of concepts*. Oxford series in cognitive development. Oxford University Press, Oxford; New York, 2009. ISBN 9780195367638 (hardcover alk. paper) 0195367634. 2008029129 Susan Carey. ill. ; 25 cm. Includes bibliographical references (p. 539-566) and indexes. Oxford series in cognitive development.
11. E. S. Spelke and K. D. Kinzler. Core knowledge. *Developmental Science*, 10(1):89–96, 2007. ISSN 1363-755X. doi: 10.1111/j.1467-7687.2007.00569.X. 118RG Times Cited:206 Cited References Count:108.
12. Annette Karmiloff-Smith. *Beyond modularity: A developmental perspective on cognitive science*. MIT press, 1995. ISBN 0262611147.
13. Steven R Quartz and Terrence J Sejnowski. The neural basis of cognitive development: A constructivist manifesto. *Behavioral and brain sciences*, 20(4):537–556, 1997. ISSN 1469-1825.
14. E Thelen and LB Smith. A dynamic systems approach to cognition and action, 1994.
15. J. Locke. An essay concerning human understanding, 1689.
16. David Hume. *An enquiry concerning human understanding*. Clarendon Press, Oxford, 1975 edn. edition, 1748.
17. Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245–258, 2017. ISSN 0896-6273.
18. Geoffrey E Hinton. How neural networks learn from experience. *Scientific American*, 267(3):144–151, 1992. ISSN 0036-8733.
19. James L McClelland and David E Rumelhart. Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114(2): 159, 1985. ISSN 1939-2222.
20. Michael J Arcaro, Peter F Schade, Justin L Vincent, Carlos R Ponce, and Margaret S Livingstone. Seeing faces is necessary for face-domain formation. *Nature neuroscience*, 20(10):1404, 2017. ISSN 1546-1726.
21. Krishna Srihasam, Justin L Vincent, and Margaret S Livingstone. Novel domain formation reveals proto-architecture in inferotemporal cortex. *Nature neuroscience*, 17(12):1776, 2014. ISSN 1546-1726.
22. N. Li and J. J. DiCarlo. Unsupervised natural experience rapidly alters invariant object representation in visual cortex. *Science*, 321(5895):1502–1507, 2008. ISSN 0036-8075. 347DW Times Cited:56 Cited References Count:46.
23. N. Li and J. J. DiCarlo. Unsupervised natural visual experience rapidly reshapes size-invariant object representation in inferior temporal cortex. *Neuron*, 67(6):1062–75, 2010. ISSN 1097-4199 (Electronic) 0896-6273 (Linking). doi: 10.1016/j.neuron.2010.08.029. Li, Nuo DiCarlo, James J 1F31EY020057/EY/NEI NIH HHS/ R01 EY014970/EY/NEI NIH HHS/ R01 EY014970-05S1/EY/NEI NIH HHS/ R01 EY014970-06A1/EY/NEI NIH HHS/ R01 EY014970-07/EY/NEI NIH HHS/ R01 EY014970/EY/NEI NIH HHS/ Neuron. 2010 Sep 23;67(6):1062-75. doi: 10.1016/j.neuron.2010.08.029.
24. Richard Held and Alan Hein. Movement-produced stimulation in the development of visually guided behavior. *Journal of comparative and physiological psychology*, 56(5):872, 1963. ISSN 0021-9940.
25. David H Hubel and Torsten N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology*, 160(1):106–154, 1962. ISSN 1469-7793.
26. RD Walk, Eleanor J Gibson, and TJ Tighe. Behavior of light-and dark-reared rats on a visual cliff. *Science*, 1957. ISSN 0036-8075.
27. J. N. Wood. Newborn chickens generate invariant object representations at the onset of visual object experience. *Proceedings of the National Academy of Sciences of the United States of America*, 110(34):14000–14005, 2013. ISSN 0027-8424. 203DU Times Cited:0 Cited References Count:47.
28. S. M. W. Wood and J. N. Wood. A chicken model for studying the emergence of invariant object recognition. *Frontiers in Neural Circuits*, 9:7, 2015. ISSN 1662-5110 (Electronic) 1662-5110 (Linking). doi: 10.3389/fncir.2015.00007.
29. J. N. Wood and S. M. W. Wood. The development of newborn object recognition in fast and slow visual worlds. *Proceedings of the Royal Society B-Biological Sciences*, 283(1829): 20160166, 2016. ISSN 0962-8452.
30. J. N. Wood and S. M. W. Wood. A smoothness constraint on the development of abstract object representations. *Cognitive Science*, 2018. doi: 10.1111/cogs.12595.
31. L. B. Smith and L. K. Slone. A developmental approach to machine learning? *Front Psychol*, 8:2124, 2017. ISSN 1664-1078 (Print) 1664-1078 (Linking). doi: 10.3389/fpsyg.2017.02124. Smith, Linda B Slone, Lauren K eng F32 HD093280/HD/NICHD NIH HHS/ T32 HD007475/HD/NICHD NIH HHS/ Switzerland 2017/12/21 06:00 Front Psychol. 2017 Dec 5;8:2124. doi: 10.3389/fpsyg.2017.02124. eCollection 2017.
32. E. D. Jarvis, O. Gunturkun, L. Bruce, A. Csillag, H. Karten, W. Kuenzel, L. Medina, G. Paxinos, D. J. Perkel, T. Shimizu, G. Striedter, J. M. Wild, G. F. Ball, J. Dugas-Ford, S. E. Durand, G. E. Hough, S. Husband, L. Kubikova, D. W. Lee, C. V. Mello, A. Powers, C. Siang, T. V. Smulders, K. Wada, S. A. White, K. Yamamoto, J. Yu, A. Reiner, and A. B. Butler. Avian brains and a new understanding of vertebrate brain evolution. *Nature Reviews: Neuroscience*, 6(2):151–9, 2005. ISSN 1471-003X (Print) 1471-003X (Linking). doi: 10.1038/nrn1606. Jarvis, Erich D Gunturkun, Onur Bruce, Laura Csillag, Andras Karten, Harvey Kuenzel, Wayne Medina, Loreta Paxinos, George Perkel, David J Shimizu, Toru Striedter, Georg Wild, J Martin Ball, Gregory F Dugas-Ford, Jennifer Durand, Sarah E Hough, Gerald E Husband, Scott Kubikova, Lubica Lee, Diane W Mello, Claudio V Powers, Alice Siang, Connie Smulders, Tom V Wada, Kazuhiro White, Stephanie A Yamamoto,

- Keiko Yu, Jing Reiner, Anton Butler, Ann B Avian Brain Nomenclature Consortium R13 MH064400-01/MH/NIMH NIH HHS/ England Nat Rev Neurosci. 2005 Feb;6(2):151-9.
33. H. J. Karten. Neocortical evolution: neuronal circuits arise independently of lamination. *Current Biology*, 23(1):R12–5, 2013. ISSN 1879-0445 (Electronic) 0960-9822 (Linking). doi: 10.1016/j.cub.2012.11.013. Karten, Harvey J England Curr Biol. 2013 Jan 7;23(1):R12-5. doi: 10.1016/j.cub.2012.11.013.
 34. Eric R Kandel. *In search of memory: The emergence of a new science of mind*. WW Norton Company, 2007. ISBN 0393070514.
 35. S. Brenner. The genetics of *caenorhabditis elegans*. *Genetics*, 77(1):71–94, 1974. ISSN 0016-6731 (Print) 0016-6731 (Linking). Brenner, S Genetics. 1974 May;77(1):71-94.
 36. H. J. Bellen, C. Tong, and H. Tsuda. 100 years of drosophila research and its impact on vertebrate neuroscience: a history lesson for the future. *Nature Reviews Neuroscience*, 11(7):514–22, 2010. ISSN 1471-0048 (Electronic) 1471-003X (Linking). doi: 10.1038/nrn2839. Bellen, Hugo J Tong, Chao Tsuda, Hiroshi T32 NS043124/NS/NINDS NIH HHS/ Howard Hughes Medical Institute/ England Nat Rev Neurosci. 2010 Jul;11(7):514-22. doi: 10.1038/nrn2839.
 37. Eleanor J Gibson and Richard D Walk. The "visual cliff". *Scientific American*, 202(4):64–71, 1960. ISSN 0036-8733.
 38. G. Vallortigara. Core knowledge of object, number, and geometry: A comparative and neural approach. *Cognitive Neuropsychology*, 29(1-2):213–236, 2012. ISSN 0264-3294. Sp. Iss. SI 012WJ Times Cited:1 Cited References Count:152.
 39. O. Rosa-Salva, L. Regolin, and G. Vallortigara. Faces are special for newly hatched chicks: evidence for inborn domain-specific mechanisms underlying spontaneous preferences for face-like stimuli. *Developmental Science*, 13(4):565–577, 2010. ISSN 1363-755X. doi: 10.1111/J.1467-7687.2009.00914.X. 611FQ Times Cited:4 Cited References Count:53.
 40. G. Vallortigara, L. Regolin, and F. Marconato. Visually inexperienced chicks exhibit spontaneous preference for biological motion patterns. *Plos Biology*, 3(7):1312–1316, 2005. ISSN 1544-9173. doi: 10.1371/journal.pbio.0030208. 949bd Times Cited:122 Cited References Count:35.
 41. E. Mascalzoni, L. Regolin, and G. Vallortigara. Innate sensitivity for self-propelled causal agency in newly hatched chicks. *Proceedings of the National Academy of Sciences of the United States of America*, 107(9):4483–4485, 2010. ISSN 0027-8424. doi: 10.1073/pnas.0908792107. 563kp Times Cited:48 Cited References Count:33.
 42. L. Regolin and G. Vallortigara. Perception of partly occluded objects by young chicks. *Perception Psychophysics*, 57(7):971–6, 1995. ISSN 0031-5117 (Print) 0031-5117 (Linking). Regolin, L Vallortigara, G Percept Psychophys. 1995 Oct;57(7):971-6.
 43. Rosa Rugani, Lucia Regolin, and Giorgio Vallortigara. Imprinted numbers: newborn chicks' sensitivity to number vs. continuous extent of objects they have been reared with. *Developmental science*, 13(5):790–797, 2010. ISSN 1467-7687.
 44. A. Martinho and A. Kacelnik. Ducklings imprint on the relational concept of "same or different". *Science*, 353(6296):286–288, 2016. ISSN 0036-8075. doi: 10.1126/science.aaf4247. Dr0em Times Cited:5 Cited References Count:26.
 45. A. Gelman and E. Loken. The statistical crisis in science. *American Scientist*, 102(6):460–465, 2014. ISSN 0003-0996. Aq9ss Times Cited:107 Cited References Count:16.
 46. Andrew Gelman. The failure of null hypothesis significance testing when studying incremental changes, and what to do about it. *Personality and Social Psychology Bulletin*, 44(1):16–23, 2018. ISSN 0146-1672.
 47. J. N. Wood. Newly hatched chicks solve the visual binding problem. *Psychological Science*, 25(7):1475–1481, 2014. ISSN 1467-9280 (Electronic) 0956-7976 (Linking). doi: 10.1177/0956797614528955. Wood, Justin N Psychol Sci. 2014 May 19;25(7):1475-1481.
 48. Stéphane Champely. *pwr: Basic functions for power analysis*, 2015.
 49. J. N. Wood. Spontaneous preference for slowly moving objects in visually naïve animals. *Open Mind*, 1(2):111–122, 2017. ISSN 2470-2986.
 50. David Peterson. The baby factory: Difficult research objects, disciplinary standards, and the production of statistical significance. *Socius*, 2:1–10, 2016. ISSN 2378-0231.
 51. Kimberly Scott, Junyi Chu, and Laura Schulz. Lookit (part 2): Assessing the viability of online developmental research, results from three case studies. *Open Mind*, 1(1):15–29, 2017. ISSN 2470-2986.
 52. Kimberly Scott and Laura Schulz. Lookit (part 1): A new online platform for developmental research. *Open Mind*, 1(1):4–14, 2017. ISSN 2470-2986.
 53. Linda B Smith, Swapnaa Jayaraman, Elizabeth Clerkin, and Chen Yu. The developing infant creates a curriculum for statistical learning. *Trends in cognitive sciences*, 2018. ISSN 1364-6613.
 54. Brandon C Roy, Michael C Frank, Philip DeCamp, Matthew Miller, and Deb Roy. Predicting the birth of a spoken word. *Proceedings of the National Academy of Sciences*, 112(41):12663–12668, 2015. ISSN 0027-8424.

Table 1. Measurement error (standard deviation) and effect size (Cohen's *d*) for all of the statistically significant conditions from 10 automated and 10 non-automated studies. Each point in Figure 1 corresponds to one entry in the table.

Authors	Year	Title	Journal	Experimental Condition	SD	Cohen's <i>d</i>
AUTOMATED STUDIES						
Wood	2013	Newborn chickens generate invariant object representations at the onset of visual object experience	PNAS	Exp 1	12%	1.24
				Exp 2	7%	2.94
				Exp 3	10%	0.86
Wood	2014	Newly hatched chicks solve the visual binding problem	Psychological Science	Color Change	12%	1.88
				Shape Change	11%	1.49
				Color-Shape Change	11%	2.39
				Binding Change	10%	1.69
Goldman & Wood	2015	An automated controlled-rearing method for studying the origins of movement recognition in newly hatched chicks	Animal Cognition	Exp 1	9%	1.08
				Exp 2-Replace one movement	3%	0.94
				Exp 2-Replace all movements	8%	1.28
				Exp 3	4%	1.78
Wood	2015	Characterizing the information content of a newly hatched chick's first visual object representation	Developmental Science	Identity Trials	15%	0.68
				Viewpoint Trials	13%	1.17
Wood & Wood	2015	A chicken model for studying the emergence of invariant object recognition	Frontiers in Neural Circuits		8%	0.88
Wood & Wood	2015	Face recognition in newly hatched chicks at the onset of vision	JEP: ALC	Edges Only	7%	5.15
				No Color	9%	3.71
				Features Only	9%	1.99
				Different Gender Coloring	17%	1.10
				Inverted	15%	1.10
				Different Age	12%	1.20
Wood, Prasad, Goldman, & Wood	2016	Enhanced learning of natural visual sequences in newborn chicks	Animal Cognition	Novel Image trials, Natural Sequence condition	9%	1.33
				Novel Image trials, Unnatural Sequence condition	6%	0.69
Wood & Wood	2016	The development of newborn object recognition in fast and slow visual worlds	Proceedings of the Royal Society B	Exp 1: Slow Object - View-invariant vs View-dependent (Within-subjects difference)	11%	0.95
				Exp 1: Fast Object - View-invariant vs View-dependent (Within-subjects difference)	16%	1.74

Wood	2017	Spontaneous preference for slowly moving objects in visually naïve animals	Open Mind	Exp 1	11%	1.57
				Exp 2: 1.25s vs 2.5s	16%	0.72
				Exp 2: 1.25s vs 5s	12%	0.74
				Exp 2: 1.25s vs 10s	11%	0.98
				Exp 2: 1.25s vs 20s	9%	1.15
Wood & Wood	2017	Measuring the speed of newborn object recognition in controlled visual worlds	Developmental Science	Exp 1: 250ms	7%	1.24
				Exp 1: 750ms	9%	0.88
				Exp 1: 1750ms	9%	0.92
				Exp 1: 3750ms	8%	1.40
				Exp 2: 125ms	5%	1.04
				Exp 2: 250ms	11%	0.72
				Exp 2: 750ms	8%	1.08
				Exp 2: 1750ms	7%	1.45
				Exp 2: 3750ms	10%	1.02
NON-AUTOMATED STUDIES						
Regolin & Vallortigara	1995	Perception of partly occluded objects by young chicks	Perception & Psychophysics	Exp 1 Complete Triangle v. Fragmented Triangle	22%	0.38
				Exp 1 Complete Triangle vs. Scrambled Triangle	34%	0.36
				Exp 1 Partly Occluded Triangle vs. Fragmented Triangle	29%	0.40
				Exp 1 Partly Occluded Triangle vs. Inversion of Occlusion	26%	0.43
				Exp 2: Reared with Occluded object, Preference for Complete Triangle vs. Fragmented Triangle	26%	0.55
				Exp 2: Reared with Fragmented Triange, Preference for Fragmented Triangle	21%	0.74
Vallortigara, Regolin, & Marconato	2005	Visually inexperienced chicks exhibit spontaneous preference for biological motion patterns	PLOS Biology	walking hen vs rigid motion	27%	0.38
				walking hen vs random motion	26%	0.38
				scrambled hen vs rigid motion	28%	0.18
				scrambled hen vs random motion	25%	0.43

				walking cat vs rigid motion	28%	0.28
				walking cat vs random motion	27%	0.27
Mascalzoni, Regolin, & Vallortigara	2010	Innate sensitivity for self-propelled causal agency in newly hatched chicks	PNAS	Exp 1	33%	0.32
Rosa-Salva, Regolin, & Vallortigara	2010	Faces are special for newly hatched chicks: evidence for inborn domain-specific mechanisms underlying spontaneous preferences for face-like stimuli	Developmental Science	Exp 1, Imprinted Chicks	38%	0.61
				Exp 1, Non-Imprinted, 1st 3 min ONLY	44%	0.25
				Exp 4	38%	0.28
Rugani, Regolin, & Vallortigara	2010	Imprinted numbers: newborn chicks' sensitivity to number vs. continuous extent of objects they have been reared with	Developmental Science	Exp 1 imprinted on 3	24%	0.28
				Exp 1 imprinted on 1	32%	0.22
				Exp 2 imprinted on 2 objects	32%	0.35
				Exp 2 imprinted on 3 objects	31%	0.30
				Exp 3 imprinted on one big object	32%	0.44
				Exp 3 imprinted on many objects	30%	0.44
				Exp 4 imprinted on heterogeneous sets	33%	0.26
Chiandetti & Vallortigara	2011	Chicks like consonant music	Psychological Science	Min 5 of testing	65%	0.24
				Min 6 of testing	60%	0.32
Regolin, Rugani, Stancher, & Vallortigara	2011	Spontaneous discrimination of possible and impossible objects by newly hatched chicks	Biology Letters	Occlusion Condition	29%	0.45
				No Occlusion	22%	0.37
				Naïve	17%	0.48
Martinho & Kacelnik	2016	Ducklings imprint on the relational concept of "same or different"	Science	Shape Condition	32%	0.39
				Colour Condition	36%	0.19
Rosa-Salva, Grassi, Lorenzi, Regolin, & Vallortigara	2016	Spontaneous preference for visual cues of animacy in naïve domestic chicks: The case of speed changes	Cognition	Exp 2	42%	0.62
				Exp 3	46%	0.44
				Exp 4	43%	0.36
				Exp 6	46%	0.29
Santolin, Rosa-Salva, Vallortigara, & Regolin	2016	Unsupervised statistical learning in newly hatched chicks	Current Biology	Exp 1	34%	0.35
				Exp 2	35%	0.35