



Enhancing Self-Regulated Learning for Information Problem Solving with Ambient Big Data Gathered by nStudy

Philip H. Winne

Today's post-secondary students are more and more frequently engaged in learning projects. Learning projects are major assignments in which students research, appraise, organize and transform information. This work typically is oriented to producing a complex and multisection document, such as a report describing a science lab experiment, a course term paper, a plan for operating a business, a course of therapy or even an honors thesis. Some learning projects are coconstructed by a team of learners. This adds complexity to the work each individual does arising from needs to coordinate people, resources, subtasks, scheduling, and sometimes shared access to resources.

Learning projects almost always call for a complex activity called information problem solving (IPS). Synthesizing Brand-Gruwel, Wopereis and Walraven's IPS-1 model (2009) and the model of self-regulated learning (SRL) I codeveloped with A. F. Hadwin (Winne and Hadwin 1998; see

P. H. Winne (✉)
Simon Fraser University, Burnaby, BC, Canada
e-mail: winne@sfu.ca

also Hadwin and Winne 2012; Winne 2013), round-trip IPS can be described in terms of six major components (cf. Eisenberg 2008; Winne et al. 2017a):

1. developing a clear understanding about resources available to support or that may impede work on the learning project;
2. designing a framework for the product of the learning project, and setting standards for judging how well information fits that framework;
3. searching for, then filtering sources of information (websites, documents) after scanning them for potential fit to the project's framework;
4. analyzing the set of filtered sources to extract information from them, and organizing these selections according to the framework previously set out;
5. planning and drafting a report, and
6. evaluating and revising the draft report to produce a final polished version satisfying requirements set for the learning project.

Learners who are productive in self-regulating learning to become better learners understand a seventh component should be added to the preceding six. That seventh component is reexamining the full scope of work to diagnose shortcomings, hypothesizing ways in which work might be improved, and planning how to launch this possibly better approach when they begin the next learning project.

Improving students' overall IPS along with component skills students engage when they work on learning projects are widely claimed to be keys to success as a student. These skills are often also described as significant contributors to personal well-being and the capacity to contribute productively to the national economy. The Social Science and Humanities Research Council of Canada implicitly takes this stance in identifying its Challenge Area #1, "new ways of learning, particularly in higher education [that] Canadians need to thrive in an evolving society and labour market" (SSHRC 2017). A survey undertaken for the Association of American Colleges and Universities was more direct in reporting "... employers indicate that they prioritize critical thinking, communication, and complex problem-solving skills over a job candidate's major field of study when making hiring decisions" (Hart Research Associates 2013, p. 4).

Post-secondary institutions recognize that they play a key role in helping students develop information literacy and other skills needed for successful

information problem solving in learning projects. Commonly, institutions approach this charge by posting guides and tutorials online, distributing handouts and posters around campus, recommending self-help books on study skills, and offering face-to-face workshops. Across these various channels, the kinds of skills addressed mainly concern annotating, how to review content, and self-test mastery of it, writing, test taking, managing time and (variously labeled) thinking critically or argumentation (e.g., Hadwin et al. 2005). A trending theme augmenting this set of basic skills is searching for information. “Google it” might be almost everyone’s first step when starting a learning project. Unfortunately, by and large, attempts to help students improve skills in all these areas are not as effective as might be expected (Hadwin and Winne 1996; Winne 2013).

One might predict learning science can remedy this situation with its wide array of research findings. In the next section, I critique a cornerstone for this belief. I argue experimental findings claimed to show “what works” are not as useful as has been claimed. This is because the preferred methodology for generating recommendations in learning science research, the randomized controlled trial (RCT), has significant limitations when one attempts to generalize experimental results as guidelines for any particular individual learner. Following my critique, I revisit and elaborate an approach to researching learning that I proposed some time ago (Winne 1992, 2006) and update in a recent article (Winne 2017a). The approach I commend rests on students’ using software to engage in online learning projects. This affords opportunity to generate big, ambient data about learning at both individual and group levels. I suggest these data can be mined to bootstrap ever more useful and empirically grounded recommendations to guide individual learners toward better learning practices and SRL that leads to improving skills for IPS.

CHALLENGES USING FINDINGS FROM RANDOMIZED CONTROLLED TRIALS TO IMPROVE IPS¹

A widely held view is that recommending changes that will benefit learning requires carrying out a “true” experiment, that is, a randomized controlled trial (RCT). Two key requirements must be met to have confidence in the findings of an RCT serving this end (What Works Clearinghouse, n.d.):

¹This section draws largely on Winne, P. H. (2017a). Leveraging big data to help each learner upgrade learning and accelerate learning science. *Teachers College Record*, 119(3), 1–24.

- (a) The sample of participants in the RCT must be representative of a well-defined population.
- (b) Participants in the RCT must be assigned to each condition investigated in the experiment by random assignment or some method that is functionally random.

Assuming other features of an RCT's methods are well done, which is no mean feat (see Shadish et al. 2002), meeting both requirements sets a stage for a recommendation like this: "If a student is a member of the same population as participants in the RCT, changing how that student goes about learning by replicating the intervention operationalized in the RCT (or in a collection of RCTs examined in a meta-analysis) has a high probability of producing a result for that student like the result that was observed for the treatment group in the RCT." I argue RCTs can build only a "fragile foundation" for strong claims like this about what works (Winne 2017a). Main points are summarized as follows.

Findings of an RCT Are Unacceptably Elastic

RCTs produce findings about differences among mean scores of groups of participants who have different experiences in the experiment. Methodologists call this difference in mean scores the effect size. It is agreed that, unless the entire population of students can participate in an experiment, the real effect size is unknowable. Because only a sample of the population of students participates in almost any experiment, the effect size observed in an RCT is just one effect size sampled randomly from a huge (theoretically, infinite) number of possible samples. This description is tacit about methodologists' conservative approach in assuming the real effect size is zero—the null hypothesis—which is tested statistically for whether data sufficiently challenge that assumption. Colloquially, what you see in one RCT is a quite fuzzy image of what you might expect to see.

Fuzziness can be illustrated quantitatively if I make a few plausible (and arguable) assumptions. First, suppose the real effect size in the population, expressed as a correlation (ρ), ranges somewhere in the range $0.10 \leq \rho \leq 0.40$. Using another common metric, Cohen's d , this range is $0.20 \leq d \leq 0.87$. Yet another expression may be clearer to readers less steeped in statistical methods. If we say the comparison group's mean score lies at the 50th percentile of all scores in the population, a treatment in this RCT would be

predicted to have a mean score ranging somewhere between the 58th percentile and the 81st percentile.

Second, no measurement is perfectly reliable. So, I assume the psychometric reliability of the outcome measure in this RCT is 0.70 in this RCT's sample.

These assumptions are backdrops for interpreting effects observed in this one RCT. Suppose it produces an effect size of $r = 0.30$, equivalent to $d = 0.63$. In this RCT, the treatment improved the mean score of students in the treatment group from the 50th percentile (the comparison group's mean score) to the 74th percentile. Now, imagine this *same* sample of students completely forgot everything about their experience in this RCT and participated in a perfect replication of the experiment. Results of this replication would have a 95% confidence interval that ranges from of $r = -0.33$ ($d = -0.70$, 24th percentile)—note: the treatment reduces achievement—to $r = 0.65$ ($d = 1.71$, 96th percentile)—the treatment is a potent benefit. It is worth pointing out the span of this interval grows *wider* to an unknown degree if *different* students experience the intervention. That, of course, must be the case when anyone applies findings of the RCT to a new sample.

As if elasticity in “the” finding of an RCT was not enough, it is actually even more elastic. A hallmark of RCTs is the C of this acronym. It identifies the trial (or experiment) as “controlled.” This means the researcher did everything possible to insure that every other factor theoretically reasoned or empirically demonstrated to affect the outcome variable was controlled, that is, did not vary. This usually makes an RCT very unlike other contexts in which the experiment's intervention might be applied. Whenever one of these factors in the real world has a value or a range different than the value it had in the RCT, the confidence interval just described blurs. It may contract. It may elongate. The center point on which it pivots may slide up or down. In short, predictions about what to expect on “replicating” the intervention studied in the RCT become quite fuzzy.

This analysis of an effect size observed in a single RCT delivers an unhappy message. It is a *very* chancy proposition to predict what to expect about the mean performance of a different group of students who experience an intervention researched in one RCT. You might counter: “A meta-analysis, where a collection of RCTs are examined, will fix this problem.” Unfortunately, no. The elasticity of the confidence interval for an effect size does not shrink if it is produced by a meta-analysis. The collection of

statistically detectable moderator variables a meta-analysis identifies help to point out which factors moderate an intervention's effects but this list does not make generalizing "the" effect any more robust.

Means Cannot Predict an Individual Student's Results

It is easy to show the mean score of a group is useless as a prediction about the score of any individual in the group, such as a student who experiences a treatment in an RCT. Statisticians model an individual's score, which they represent symbolically as X_{ij} , in terms of three components. One component, symbolized by μ , represents the mean score of all students in the population to which an individual student belongs. The second component, τ_j , indicates how much the population mean μ changes as a result of everyone in the population experiencing the treatment. The subscript j is appended to signal a particular group in the experiment—say, the treatment group (1) or the comparison group (0)—to which a particular student belongs. Finally, the third component is an "error" term, symbolized ε_{ij} . This component reflects how much the score of a particular individual i who is in group j differs from the mean score of all the other students from the population who are in group j . Putting these all together, the i th individual student who is group j has a score equal to the sum of these three components. Eq. (8.1) shows how this is represented when a weighting factor, b , is applied to the term representing the effect of a treatment. If the student is in the treatment group, $b = 1$. If the student is in the comparison group, $b = 0$ and the effect of the treatment (τ_j) is nullified when it is multiplied by zero.

$$X_{ij} = \mu + b_j \tau_j + \varepsilon_{ij} \quad (8.1)$$

What is critical to know about how statisticians use this expression when they construct a statistical analysis is this: When an average score in a group, say the j th one, is calculated by aggregating all the individual students' scores, it is assumed the sum of all the individual ε_{ij} components will be zero. The average obliterates individuality; the ε_{ij} term is zero. If this is not the case, the analysis suffers bias. In short, knowing the mean score of a group of students offers no help in predicting what the score will be for any individual student.

Populations Are Ineffectually Described

At the outset, I noted the validity of inferring the presence of an effect in an RCT depends on the sample of participants in the experiment being a *random* sample from a *well-defined* population. The requirement that the sample be random arises because this is a critical assumption underlying statistical models used to investigate inferences about whether an intervention in an RCT produces a statistically detectable effect. If the sample is not randomly drawn from its population, inaccuracy is introduced into the inference of whether an effect appeared in the experiment. I believe it safe to say the vast majority of RCTs fail this requirement. Samples participating in RCTs are almost always samples of convenience. I (Winne 2017a) coined the term “pseudo-random controlled trial” or P-RCT to reflect this situation.

Another issue arises regarding the requirement that the population be well defined. In this context, “well defined” means the population is defined by factors that are empirically and reliably known to cause variation in the outcome variable measured in a P-RCT. Why is this a requirement? Consider an experiment in which high school students in the intervention group study new terms that will be important in an upcoming lesson. They study until every student can perfectly define each term from memory. Peers in the comparison group join these students after the intervention group’s preparatory session, and all students are shown a movie about farming practices before they all take exactly the same achievement test. Now, suppose I define the population as 14-year olds and 60% female. (The *APA Publication Manual* requires noting these two demographic features of participants in studies.) My sample of 30 students in the intervention group, being a randomly lucky one, consists of only 14-year-olds and has 18 (60% of 30) girls.

Does age cause variation in, say, understanding why the early stages of root infection are greatly affected by soil pH? No, age is a poor proxy for opportunity to learn about content. Moreover, age is not at all a useful proxy for the quality of those opportunities as learning experiences. Is sex a cause of learning or lack thereof? No, it is a poor proxy for potential interest in or opportunity to engage in learning about this topic. Because these “defining” factors are poor proxies, and other factors that really cause variation in the outcome measure in this experiment are unspecified, this study suffers what statisticians call the specification error. If an effect is detected, there is neither theoretical nor empirical warrant to generalize

the effect to people who happen to be 14 years old or have any basis for considering how much the effect might vary if the population of students is only 52% female or all female.

Summary

The randomized controlled trial (RCT), much acclaimed in research on learning science, is really a pseudo-randomized controlled trial (P-RCT). Even if shortcomings of P-RCTs could be redressed, a probabilistically inferred effect is limited to describing differences between the mean scores of groups. It cannot forecast what to expect for any particular student who later experiences the same treatment studied in the experiment. If we hold a view that each student is responsible for carrying out operations that generate learning, P-RCTs have significant limitations as sources of robust recommendations about “what works for you.”

I hasten to emphasize the bulk of learning science should not, *ipso facto*, be ignored or discarded. Later, I suggest an important role for the current body of research in learning science. Moreover, I argue learning science can be accelerated and lend much value to helping students tackle learning projects and other IPS tasks.

To reach this point, I first set out goals to be achieved by a new approach to research. Then, I explore how modern software systems can support approaching those goals. The prize at the end of this journey is a systematic plan for helping individual students productively tackle information problem solving in the age of nearly unlimited online resources.

GOALS FOR RESEARCH ON IPS

There are two key goals for a new approach to research on promoting IPS. One is to describe, in terms that each student can understand and act on, what should replace the tactics and strategies they presently use in IPS that do not produce optimal results. It is important to emphasize: *Students* need to be clear about an intervention’s operational definition because *they* enact the “treatment.” Related to this goal are two assumptions.

Learning science commonly identifies the ways an author can structure information presentations (texts, videos, diagrams, etc.) to improve IPS. However, I assume very few authors posting information to the Internet know about or care to implement these recommendations. If my conjecture is valid, students must “fend for themselves” when they source and analyze information for IPS.

Principles of learning science are typically expressed in terms of unobserved constructs, e.g., rehearsal, decay, metacognitive monitoring, elaboration, etc. To most students, these are foggy notions. In contrast, students could readily understand operational definitions of these constructs, for example, considering the foregoing constructs: Retype text you highlighted. Of text you highlighted, you did not recall this (particular information.) List standards you use to decide whether you recall enough about this (particular information). Illustrate this principle by an example from your experience.

I recommend research in IPS be designed to generate learning analytics expressed in terms of operational definitions rather than theoretical constructs. I predict when learning analytics have this form, students can understand clearly what they did when they studied previously and what they can consider as supplements to or replacements for prior actions to improve learning.

A second goal for the new approach to research I recommend for improving students' IPS is to trace, as much as possible, everything students do *as* they work on IPS tasks. If this goal can be achieved, RCT's fetters to random sampling and random assignment can be cut. Here is why.

In the classical approach of experimenting to identify an intervention's effect(s) on outcomes, it is axiomatic numerous unknown factors causally affect the outcome. This is, in fact, how the normal distribution of scores acquires its shape. When the number of participants is large enough, and when other experimental controls are sufficiently well implemented, random selection and random assignment of participants to an experiment's conditions provide mathematical insurance that causes with positive influence and with negative influence "balance out." When this is the case, interpretations about the effect(s) of an intervention do not suffer confounds such that an effect is just as likely attributed to the intervention rather as to some unknown causal factor(s).

There are problems with this. Randomly selecting a sample of participants from a well-defined population is almost never achieved. Setting this aside, sample sizes would typically need to be prohibitively large. For example, a meta-analysis by Bakermans-Kranenburg et al. (2005) nominated six factors as moderator variables affecting the outcomes of early childhood interventions. If a new RCT was planned to investigate a new intervention, I calculated the size of a random sample of children from a well-defined population would need to be approximately 12,960 to avoid a confounded interpretation about whether the treatment was beneficial (or harmful) (Winne 2006). While samples of that size might occasionally

be achieved, interventions requiring special training or unconventional environmental designs suffer erosion of control when samples are so large.

To recap, new research on IPS should strive to provide students with learning analytics describing operationally how to study rather than constructs that matter theoretically. To generate such learning analytics requires data fully describing what each student does while studying in every studying episode. Next, I describe software designed to accomplish meet this standard.

nSTUDY: SOFTWARE FOR EVERYDAY IPS THAT GENERATES AMBIENT BIG DATA

nStudy is an extension programmed for the popular Google Chrome web browser. nStudy's features are tools learners can use to operate on information presented in web pages, pdf documents and videos they find on the Internet. The software was designed, in part, to open the black box of learning, that is, to bring into observable form learners' cognitive operations on information and motivational states that shape what they learn (Winne 1982; Winne et al. 2017b).

Suppose Noah's project is to argue whether owning a hybrid car is a wise consumer choice. After opening his browser and logging into nStudy, Noah enters "hybrid cars pros and cons" in nStudy's search box. From Google's returns, he selects one source he judges should be rather positive (www.pluginincars.com) and begins reading. Early in the text, he drags his cursor over "the cost per mile to fuel an EV is approximately one-third to one-quarter the cost of gasoline (on a cost per mile basis)" (Berman 2016). As soon as he lifts his finger from the trackpad, nStudy pops up a menu of options for operating on that selected information: *quote*, *note*, *term* (Fig. 8.1). Noah chooses *quote* and creates a tag for the information he selected, "pro." In response, nStudy (a) highlights the text Noah selected, (b) paints a small colored nub next to the scrollbar (a region we call the gutter) to mark the quote's relative location in the web page so Noah can see where he's made quotes, (c) adds beneath the Tags header in of the sidebar the new tag Noah created, and (d) copies the text Noah selected to a sector of nStudy's sidebar as a quote (Fig. 8.2).

Table 8.1 shows data nStudy records about Noah's work so far in a database on a server. Every event is time-stamped, accurate to at least 1/100th second. These data are ambient data; they are "collected as a matter of course"

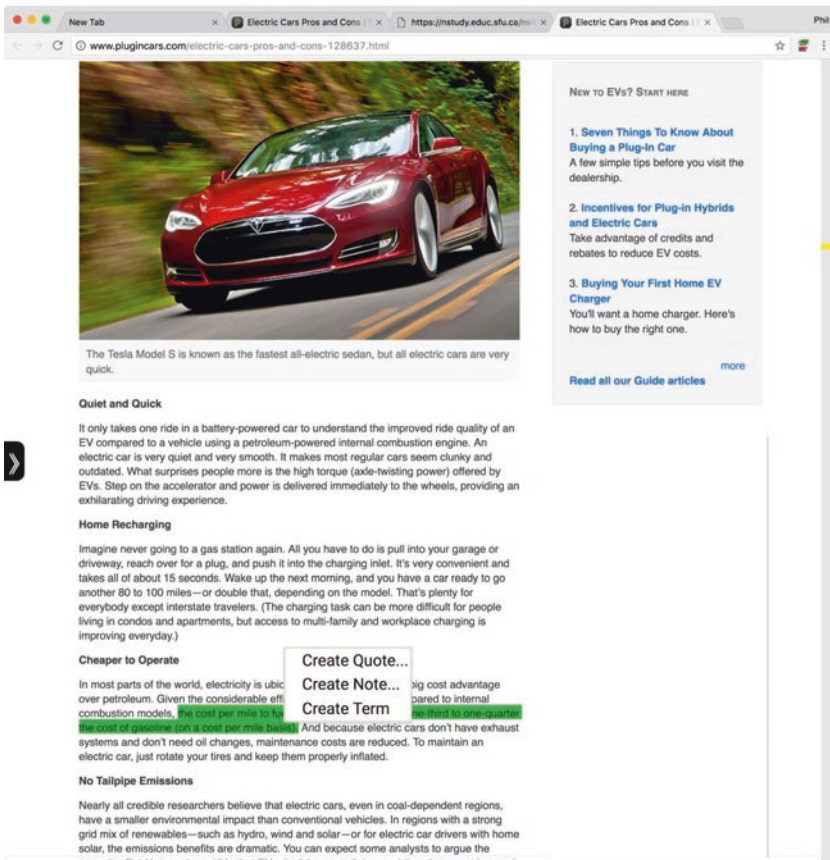


Fig. 8.1 Noah selects information to quote and tag “pro”

(Pistilli et al. 2014, p. 85) and trace, as much as possible, every event that could be observed as Noah worked on his learning project.

The principle guiding nStudy’s design was to record everything possibly observable about learners’ studying activities so that a full account of the studying episode was created. nStudy’s trace data help open the black box of learning (Winne 1982). For example, Noah’s choice of search terms reveals his judgment about key components in crafting an argument: pro, con. His search query also reveals a choice to focus on, or ignorance about, the difference between hybrids, all-electric plug-in cars,

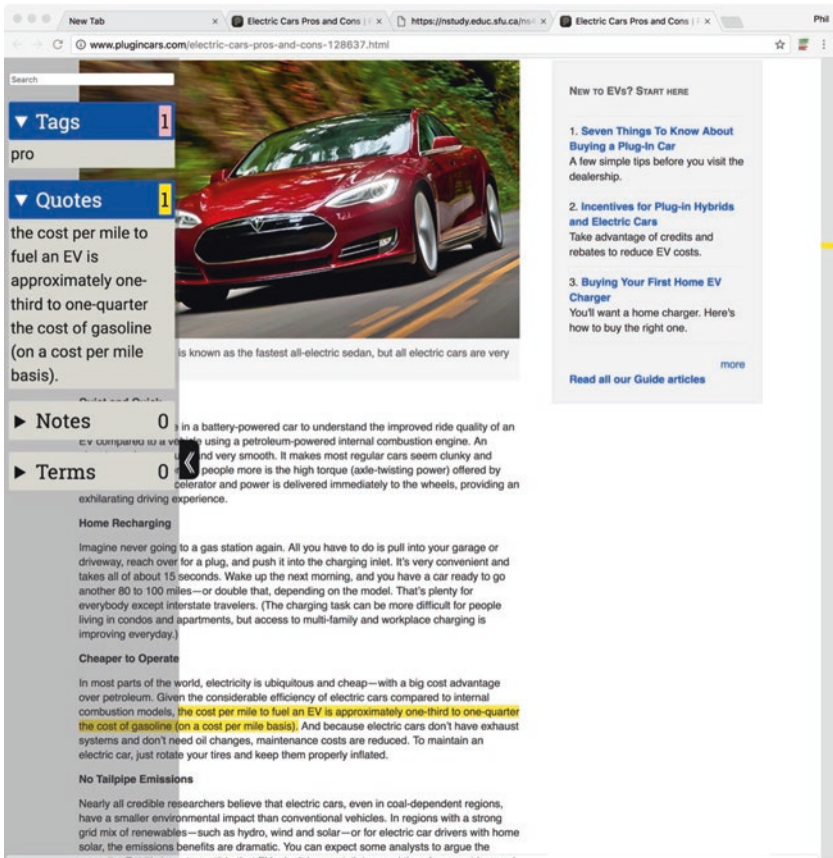


Fig. 8.2 nStudy's trace data mirror Noah's operations on information quoted and tagged

and hybrid plug-in cars. Noah's selection of text about cost per mile of ownership signals he was metacognitively monitoring information in the web page. The selection he quoted satisfied standards used in metacognitive monitoring. His specific standard is revealed by the tag he created.

nStudy offers an array of other features that students can use in IPS tasks. These features afford gathering data that can recreate the temporal, operational, and informational aspects of work on learning projects. Brief descriptions of artifacts and features in nStudy are provided in Table 8.2.

Table 8.1 Operations and information operated on in Noah's studying session

<i>Operation</i>	<i>Information</i>
Log in	Identity
Choose search	Place cursor in search box
Search	Search terms entered: hybrid cars pros cons
View	Search results
Choose URL	www.pluginincars.com
View	Content of the page at www.pluginincars.com
Choose text	The cost per mile to fuel an EV is approximately one-third to one-quarter the cost of gasoline (on a cost per mile basis)
Choose operation	Quote option on popup menu
Create tag	"Pro"
Quote	The cost per mile to fuel an EV is approximately one-third to one-quarter the cost of gasoline (on a cost per mile basis)

When learners create, edit, file, review, or destroy an artifact, nStudy logs the complete interaction. Similarly, when learners interact with a feature, for example, searching for an artifact, that interaction is also logged.

nStudy is now being extended to retrieve and organize data from one or multiple learners' databases for input to computations generating learning analytics. In addition to straightforward mirror reports (e.g., "You made 18 notes in today's 2 studying sessions, 8 of which you tagged *review later*.) comparisons can be ipsative (within one learner across time), criterion referenced or norm referenced (relative to a defined group of a learner's peers). Because inputs to learning analytics are operations a learner applies using nStudy's tools, for example, quoting or searching for a note about a particular topic, and because the information a learner operates on is recorded in nStudy's database, learning analytics can be presented in terms learners are accustomed to using when they use the software.

How Software Helps

Data that nStudy gathers approach big data. A common 500-page introductory textbook containing approximately 1000 terms, wherein the learner makes 2 quotes per page and 2 notes per page, plus searches for and reviews 25% of these artifacts, will generate approximately $500 \times 2 \times 2 + 1000 + .25 \times 3000 = 4000\text{--}5000$ records per student-semester. An estimate of data points generated in a typical IPS project, such as a term

Table 8.2 nStudy's feature set

<i>Artifact/ feature</i>	<i>Description</i>
Bookmark	Titled, clickable and searchable links to source content (pages, pdfs) in the Internet
Quote	Text selected in a source that is copied to and searchable in a learner's workspace. Clicking a quote returns the learner to its source and the quote in context
Note	Learners' elaborations and interpretations of quotes. Replacement text in a text field guides learners' entries, e.g., "title your note" or "describe your feeling"
Note form	Schemas guiding analysis of content. Forms can be configured with multiple kinds of fields: field labels, text, checkboxes, radio buttons, sliders, drop-down lists, dates, images, attached files and "see also" field which records hyperlinks between the active note and other nStudy artifacts
Term	Notes that name and describe key terms, characters, events, etc.; the term form includes a "see also" field
Termnet	A node-edge network in which nodes are terms and edges signify a relation. Clicking a node opens the term note for quick review. When a bookmark, essay or discussion (see Hub) is opened, nStudy identifies its subgraph of terms—a local glossary.
Tag	An index to be applied to a set of nStudy artifacts to classify them conceptually (e.g., "conjectures"), motivationally (e.g., "interesting"), qualitatively (e.g., "vague"), or in relation to tasks (e.g., "review for test). Tags can be restricted to a set provided by a researcher/instructor or freely created by a learner
Essay	A feature for drafting and formatting compositions, lab reports, business plans, diagnoses, etc. Quotes, notes, and terms can be incorporated into an essay by dragging or copy/pasting them into an essay
Hub	A feature for communicating with peers synchronously or asynchronously. Any nStudy artifact (bookmark, quote, note, term, essay, map) can be sent to collaborators
Library	A feature for browsing and searching for artifacts by title, content, kind, temporal attributes (e.g., edited this week) and other artifact metadata
Map	A feature for constructing and displaying relations (links) among nStudy artifacts (nodes); a concept map. Learners construct maps by adding items and linking them at will; or by filtering the library to a desired subset, then clicking a button, "map it." Artifacts in maps can be grouped to form submaps. An item's "conceptual neighborhood" can be shown at arbitrary "distance" measured by the number of links traversed
Queries	A special folder in the library of researcher-provided or learner-designed search expressions with optional fill-in text fields, e.g., "find notes made last week not reviewed about [topic]"

paper, might be 500. If a typical learner enrolls in four courses and data are available for a freshman class of 5000, the flow of data is approximately 100 million raw data points per semester, not counting time stamps and semantic features of information on which learners operate. Because nStudy's data share a common format independently of where students enroll, post-secondary institutions pooling data would expand this volume. There can be big data about how learners study.

Ambient data of this volume gathered in the natural ecology of studying and IPS projects offer significant affordances for mining to trace how students study and orchestrate their work in IPS projects. These data also can be mined to identify how naturally arising events as well as inserted learning analytics perturb patterns of IPS and nudge achievement (Winne and Baker 2013). It is practically inevitable within this volume of data that some naturally occurring patterns will validate findings already developed in learning science. It is also highly likely new patterns will be discovered that have not yet been explored in learning science. A particularly appealing opportunity is the ability to track the nature of IPS skills, study tactics, and their adaptations arising from injections of learning analytics over time. While time spans fall considerably short of "life long learning," there is an exceptional prospect of mapping the developmental trajectory of IPS skills over several years of an undergraduate career.

An additional advantage of software systems like nStudy is future capability to offer just-in-time responsive support for students. In conventional research programs, data for a single experiment may be gathered over 1–2 months. This is followed by a period where data are analyzed, a paper is drafted and submitted for publication, the publication is published with a lag of 1–2 years and, perhaps 2–5 years after that, someone synthesizes multiple studies in a meta-analysis. In contrast, when ambient data are collected using software systems like nStudy, data are immediately available for automated or hand-crafted analyses. The concept of a contained study disappears. It can be replaced by "overnight" updates to the state of the art. As and when findings cohere, so recommendations can be justified, these can be distributed directly and immediately to students. This sparks a continuous and rapidly responsive cycle of evidence-based investigation and adaptation. In addition, students are relieved of having to visit a tutor or study skills center. They can receive advice tailored to their idiosyncratic approach to IPS as they log in to the next study session.

CONCLUSION

Online information is a major resource students mine in learning projects, less extensive IPS tasks, and everyday academic work. With robust regard for privacy, as learners use online technologies for studying and other IPS tasks, unparalleled prospects arise to gather extensive ambient data about every learner's work over time and across learning projects. The data generated are significantly greater in volume and detail than has been possible to realize in a pen-and-paper world. But volume is actually a drawback unless it can be intelligently mined and analyzed. Powerful tools for mining and analyzing large sets of data knock down this potential barrier.

Data that software systems can gather will accelerate learning science and enhance learners' achievements (Winne 2006, 2017a, b; Winne et al. 2017b). The learning management systems in widespread use across post-secondary institutions miss this opportunity because data they gather cannot reveal what learners *do* in the IPS "activity stream" (DiCerbo and Behrens 2014).

As previously forecast, prior findings and models developed through diligent work in prior learning science should not be set aside in a frenzy to gather big ambient data about learning. Current findings are the best available hypotheses about how to support learning, increase motivation and enhance achievement. The very attractive opportunity afforded by software systems like nStudy is a hugely increased capacity to test such findings further and more penetratingly. Opportunity to identify moderator and mediator variables is hugely advantaged in this technologically supported ecology. And, because data are big, the pool of data can be stratified to finer grain and arrayed in more complex combinations than today's experimental samples or meta-analyses allow. Moreover, the scope of ambient trace data made available by online systems like nStudy opens doors to explore new research questions in far less costly and much greater variety than is possible today. A new era in research in IPS supported by software systems may help learning science evolve to its next level.

REFERENCES

- Bakermans-Kranenburg, M. J., van Ijzendoorn, M. H., & Bradley, R. H. (2005). Those who have, receive: The Matthew effect in early childhood intervention in the home environment. *Review of Educational Research*, 75, 1–26.
- Berman, B. (2016, October 14). *Electric cars pros and cons*. Retrieved January 8, 2016, from <http://www.plugincars.com/electric-cars-pros-and-cons-128637.html>

- DiCerbo, K. E., & Behrens, J. T. (2014). *Impacts of the digital ocean on education*. London: Pearson.
- Eisenberg, M. B. (2008). Information literacy: Essential skills for the information age. *Journal of Library & Information Technology*, 28(2), 39–47.
- Hadwin, A. F., & Winne, P. H. (1996). Study skills have meager support: A review of recent research on study skills in higher education. *Journal of Higher Education*, 67, 692–715.
- Hadwin, A. F., & Winne, P. H. (2012). Promoting learning skills in undergraduate students. In M. J. Lawson & J. R. Kirby (Eds.), *Enhancing the quality of learning: Dispositions, instruction, and mental structures* (pp. 201–227). New York: Cambridge University Press.
- Hadwin, A. F., Tevaarwerk, K. L., & Ross, S. (2005, April). *Do study skills texts foster self-regulated learning: A content analysis*. Paper presented at the annual meeting of the American Educational Research Association, Montreal, Quebec.
- Hart Research Associates. (2013). *It takes more than a major: Employer priorities for college learning and student success*. Washington, DC: Author. Retrieved November 30, 2015, from https://www.aacu.org/sites/default/files/files/LEAP/2013_EmployerSurvey.pdf.
- Pistilli, M. D., Willis, J. E., & Campbell, J. P. (2014). Analytics through an institutional lens: Definition, theory, design, and impact. In J. A. Larusson & B. White (Eds.), *Learning analytics: From research to practice* (pp. 79–102). New York: Springer.
- Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston: Houghton Mifflin.
- SSHRC. (2017, February 21). *Future challenge areas*. Retrieved from http://www.sshrc-crsh.gc.ca/funding-financement/programmes-programmes/challenge_areas-domaines_des_defis/index-eng.aspx
- What Works Clearinghouse. (n.d.). *Procedures and standards handbook, version 3.0*. Retrieved from http://ies.ed.gov/ncee/wwc/pdf/reference_resources/wwc_procedures_v3_0_standards_handbook.pdf
- Winne, P. H. (1982). Minimizing the black box problem to enhance the validity of theories about instructional effects. *Instructional Science*, 11, 13–28.
- Winne, P. H. (1992). State-of-the-art instructional computing systems that afford instruction and bootstrap research. In M. Jones & P. H. Winne (Eds.), *Adaptive learning environments: Foundations and frontiers* (pp. 349–380). Berlin: Springer.
- Winne, P. H. (2006). How software technologies can improve research on learning and bolster school reform. *Educational Psychologist*, 41, 5–17.
- Winne, P. H. (2013). Learning strategies, study skills and self-regulated learning in postsecondary education. In M. B. Paulsen (Ed.), *Higher education: Handbook of theory and research* (Vol. 28, pp. 377–403). Dordrecht: Springer.

- Winne, P. H. (2017a). Leveraging big data to help each learner upgrade learning and accelerate learning science. *Teachers College Record*, 119(3), 1–24.
- Winne, P. H. (2017b). Learning analytics for self-regulated learning. In C. Lang, G. Siemens, A. Wise, & D. Gašević (Eds.), *Handbook of learning analytics* (pp. 241–249). Beaumont: Society for Learning Analytics.
- Winne, P. H., & Baker, R. S. J. D. (2013). The potentials of educational data mining for researching metacognition, motivation and self-regulated learning. *Journal of Educational Data Mining*, 5(1), 1–8.
- Winne, P. H., & Hadwin, A. F. (1998). Studying as self-regulated learning. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Metacognition in educational theory and practice* (pp. 277–304). Mahwah: Lawrence Erlbaum Associates.
- Winne, P. H., Vytasek, J. M., Patzak, A., Rakovic, M., Marzouk, Z., Pakdaman-Savoji, A., Ram, I., Samadi, D., Lin, M. P. C., Liu, A., Liaqat, A., Nashaat, N., Mozaffari, Z., Stewart-Alonso, J., & Nesbit, J. C. (2017a). Designs for learning analytics to support information problem solving. In J. Buder & F. W. Hesse (Eds.), *Informational environments: Effects of use, effective designs* (pp. 249–272). New York: Springer.
- Winne, P. H., Nesbit, J. C., & Popowich, F. (2017b). nStudy: A system for researching information problem solving. *Technology, Knowledge and Learning*, 22(3), 369–376.