Addiction Paper Authoring Tool (PAT): A Guide

Robert West[1]

1 University College London, University of London

Abstract

The Addiction Paper Authoring Tool (PAT) is an online application designed to improve the writing of research reports in the field of addiction. It is currently restricted to 2-arm randomised controlled trials but will be extended to cover other major types of study design. It prompts authors for information required, based on an underlying ontology of the field, and provides guidance and structure to ensure that the information provided is clear, consistent, comprehensive and coherent. It outputs a Word version of the paper that can then be edited for style. Equally importantly, it outputs a machine-readable record that can provide a basis for automated searching and evidence synthesis.

What is Addiction PAT?

The Addiction Paper Authoring Tool (PAT) is an online application for writing research papers in the field of addiction. The current version is limited to reporting 2-arm randomised controlled trials, but it is planned to extend it to handle other major categories of study design.

How can I use PAT?

Simply go to addictionpat.org, register as a user and start to write a paper. You will need an ORCID (orcid.org) You can invite other authors by getting them to register and then including them as an author via their email address. Their details will be automatically imported and they will be able to see the paper in their list of papers. Clicking on the paper will allow them to view and edit it.

Using PAT collaboratively

Information entered into PAT is saved automatically as you go. More than one author can amend a paper at the same time.

How did the need for PAT emerge?

The need for PAT was highlighted by:

1.   evaluating research reports being submitted to the journal, Addiction, which without exception were missing crucial items of information on submission, and many of which were still missing important information after having completed the review process.

2.   evaluating published research reports in other major journals, including the leading medical journals, and finding that in many cases even these were missing key items of information needed to evaluate or interpret the findings.

3.   attempting to manually extract detailed information from published research reports to be able to synthesise evidence on particular research questions and finding that much of the information required was missing completely, very time-consuming to find or expressed in a way that was unclear and not comparable with other reports.

4.   attempting to develop automated tools for extracting information from research reports and finding that missing information, a high degree of unnecessary variability and ambiguity in the way information is conveyed, made this impracticable to achieve with a high degree of accuracy.

5.   contributing to writing of research reports and being aware of the high degree of inefficiency in the process with very large numbers of drafts being required even before the point of submission to a journal.

These problems are not limited to addiction research and are still prevalent despite the development and use of reporting guidelines.

How does PAT address these problems?

With PAT, authors are provided with detailed online prompts for all the information required in their report. The prompts are tailored based on the information already provided. Authors are also guided as closely as possible in the way that they provide the information, with the use of fixed response categories where appropriate but with the opportunity to write free text where needed. The combination of detailed, intelligent prompts, and a combination of fixed categories and free text improves the chances that the information provided is complete, consistent and expressed in a way that is clear and comparable with other research reports.

For some parts of the paper, PAT builds the text in the paper as the information is entered so that users can see, as they type, what the text will look like. A particularly important example is the 'Research Question Wizard' which prompts the user for information and constructs a well-formed research question that is precise and can form

the basis for structuring the results and conclusions.

For example the question 'What is the effectiveness of the UK NCSCT standard treatment programme compared with usual care in achieving an increase in the percentage achieving 6-month continuous abstinence in smokers who attend stop smoking services in Malaysia?' is based on the user choosing 'effectiveness study' rather than 'efficacy study', the short label for the intervention is 'UK NCSCT standard treatment programme', the short label for the comparator is 'usual care', the outcome is 'percentage achieving 6-month continuous abstinence', the population is 'smokers who attend stop smoking services' and the setting is 'Malaysia'.

What are the outputs from PAT?

At any time in the paper-writing process, authors can preview on screen a draft of the paper being produced. They can also generate a Word file that can be edited as required for style and to meet the requirements of a given journal.

In addition to the Word file, PAT generates a data record that contains all the information entered in machine-readable form. This can be stored as a supplementary file when the article is published to facilitate automated data extraction for evidence synthesis.

PAT does not provide a fully stylistically perfect manuscript. Its goal is to ensure that all the necessary information is presented but stylistic issues will still have to be addressed during editing of the Word file once the draft version has been produced.

How has PAT been developed?

The fields to include in PAT records were identified from:

1. Key author guidance documents (e.g. CONSORT and its extensions - http://www.consort-statement.org/) and an online protocol-writing tool (https://e-protocol.od.nih.gov/#/home).

2. Information identified as needing to be extracted from RCTs in the Human Behaviour Change Project (https://www.humanbehaviourchange.org/), a major collaborative project to develop an AI system to use natural language processing to read behaviour change evaluation reports and make outcome predictions based on information about interventions, the target population and the setting.

3. Detailed examination of published RCTs in the area of addiction to identify all the information that needs to be presented.

4. Development of a set of fields from a complex RCT

(https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6629841/) that was being written as a test case.

An iterative process was used to develop a user interface that would allow flexibility for authors to enter information in a way that was natural for them while ensuring that, where there was a logical ordering of provision of information, this was followed.

A version of PAT was tested and revisions undertaken as required until a minimum viable product was produced. Since then the process has been ongoing as with any software product of this kind.

What is the structure of PAT?

PAT consists of an author registration process plus 33 'wizards' for providing information for a given paper.

Author registration: The author registration module allows users to set up and modify accounts and provide information that will be used in any paper that is written, e.g. institutional address, ORCID, declaration of competing interest. This information is imported into all papers and can be updated as required.

For each paper the wizards are modules that prompt users for information of various types
.

Author: This allows a user to add authors (who are registered) so that they can collaborate on the paper, state their contribution using the CRediT taxonomy (https://casrai.org/credit/), and provide other paper-specific information.

Paper metadata: This prompts users for information about the paper including the title, short-title, key words, supporting organisations and funding, other acknowledgements, study registration information and ethics approval information. Where relevant structured information is requested such as the type of support provided and a classification of the type of supporting organisation (e.g. government, industry, charity). Having this information in structured, machine-readable form will facilitate classification of papers and analyses to assess risk of bias.

Abstract: This prompts users for information for the abstract in a detailed and structured way that helps to ensure that the abstract is complete and accurate. For

example, it requires information about the background and aims, details of the intervention and comparator groups, the numbers recruited to each, the setting, key participant characteristics, outcome measures, findings, and conclusions.

Study topics and aims: This wizard creates the first paragraph of the Introduction. It outlines the topic of the study, states the problem that the study seeks to address and specifies the focus of the study, including the target behaviour or any health outcomes. It provides a frame and templates for the type of language to use. It also creates a single paragraph stating the aims of the study and prompts users for particular kinds of language with which to express those aims.

Research questions: This wizard creates the final paragraph of the Introduction. It populates fields each of which contains a research question. It prompts users for a frame for the question from a defined list and then for terms to be entered to complete the question. It builds the research question as each item of information is entered so that users can check that it meets their requirements. It results in a well-formed, detailed research question that is machine-readable and used as a basis for structuring other parts of the paper.

Importance and timeliness: This wizard populates a field containing paragraphs in the Introduction that describe the rationale for the topic of the study in relation to each research question. It provides a frame for the paragraphs and prompts for information to be entered into subfields. This wizard is linked to the Research Question wizard to ensure that all the research questions are fully justified.

Reason for study design: This wizard populates a field containing paragraphs in the Introduction describing the rationale for the specific study design, including the choice of measures and choice of comparator. It provides a frame and prompts for information to be entered into subfields. This ensures that authors explain their reasoning for the choices made in the study to help with the interpretation of the results. This information is often missing or incomplete in study write-ups.

Mechanism of action: This wizard populates a field containing paragraphs in the Introduction describing the logic model (the mechanism by which the intervention is expected to have an effect relative to a comparator). It provides a frame for describing the model and the evidence and analysis supporting the model. Often the logic model for interventions is missing, incomplete or presented in a way that is unevaluable. The

specific prompts in this wizard help to mitigate this. This is a field that will be developed further using an ontology-based model-building system that is being developed (https://europepmc.org/article/med/30962614).

Design: This wizard creates a single paragraph describing the study design in the Methods section. It populates this from subfields that select from fixed lists of study designs and qualifiers. This is a part of the report that authors often struggle with. The wizard provides detailed prompts to build this paragraph ensuring consistent use of terminology and phrasing, including the type of comparator group, type of randomisation, method of randomisation, randomisation ratio etc.

Setting: This wizard creates a single paragraph describing the study setting in the Methods under the 'Setting' subheading. Setting is often poorly described in study reports; this wizard provides detailed prompts to mitigate this and in due course will use an ontology of intervention settings to ensure that these are described in ways that will be machine-readable and comparable across studies.

Participants: This wizard creates a single paragraph in the Methods to describe the study participants (individuals, groups, sub-populations or populations), how they were recruited, the exclusion and inclusion criteria used to select participants, and the reason for those criteria. Using this wizard should reduce omissions and ensure that participant information is presented in a way that is consistent across studies.

Sample size determination: This wizard creates a field containing a single paragraph in the Methods section providing key information about the basis for the sample size determination and detailed of calculations undertaken, including the statistical test, alpha level and expected effect size.

Procedures: This wizard creates a section in the Methods describing the sequence of events in the study and information about those events. This wizard provides a sophisticated user interface that presents a list of possible events that may occur during a given study and allows authors to drag them into the order in which those events occurred. It ensures that crucial information is made clear, such as whether consenting occurred before or after randomisation. This will radically improve the clarity of reporting and also allow automated information extraction of this information from study reports.

Intervention and comparator: This is a sophisticated wizard that implements a more

detailed version of the TIDieR framework (https://www.equator-network.org/reporting-guidelines/tidier/) for intervention reporting. It breaks down interventions into components parts. It currently allows users to characterise the content of interventions using an established taxonomy of behaviour change techniques (https://academic.oup.com/abm/article/46/1/81/4563254). It prompts for information about the mode of delivery and details of the delivery source. It presents the intervention and comparator side-by-side so that users can copy across components and make clear precisely how the intervention differs from the comparator.

Constructs and measures: This wizard populates a field containing paragraphs in the Methods section, describing the constructs measured in the study and the ways that these are measured. This wizard is used to pre-populate tables of baseline characteristics and outcomes to ensure complete concordance and no loose ends (which are common in RCT reports). It makes a distinction between the construct being measured (e.g. cigarette dependence) and the measure (e.g. Fagerstrom Test for Cigarette Dependence) to avoid the common problem of conflation of a measure of something with the thing itself, and to make it easy for automated review systems to map constructs on to measures. It has an intelligent prompt system that uses information about the type of measure (e.g. categorical) to prompt users when reporting results of the measure to provide the appropriate kind of data (e.g. proportion). It also prompts for information about the validity or reliability of measures, information that is mostly lacking in RCT reports but it important for interpreting findings. Finally it requires users to classify measures in terms of 'primary outcome', 'secondary outcome' (including 'process'), and 'baseline'.

Analysis: This wizard populates a paragraph in the Methods section precisely specifying the statistical analyses undertaken. It has detailed prompts to ensure that users provide all the information required, including details of how missing data were handled. Given the increasing importance of calculating Bayes factors for hypothesis testing it prompts for information required to interpret these. It repeats this for all the research questions.

Deviation from protocol: This wizard provides detailed prompts for users to communicate how if at all the study deviated from the pre-registered protocol. The structured format and level of detail should enable a much greater level of understanding to be gained from the study report than is typically the case. It should also help with the review process, reducing the need for reviewers to scour protocols and study reports for differences.

Baseline characteristics table: This wizard populates a field containing a table in the Results section. It describes participants' baseline characteristics in the total sample and in each group. It prompts the user for information concerning all the baseline measures and ensures that the information is presented in a consistent, machine-readable form as well as producing a Word table for the human-readable output.

Baseline characteristics text: This wizard is a single free-text field that allows users to describe the baseline characteristics of the study sample, highlighting key information from the table.

Main results table: This wizard populates a field containing a table in the Results section describing results in relation to the primary research question. It can also include results related to a secondary outcomes providing that these do not include sub-group analyses. As with the baseline results table, it uses information from the 'Constructs and measures' wizard to prompt for the appropriate information. It presents the findings in a standardised way that helps ensure comparability across studies and is machine-readable. It turns out that extracting results information from current RCT reports is one of the hardest NLP tasks and is also very difficult and time-consuming for human reviewers. PAT should greatly improve the efficiency of this process.

Mains results text: This wizard allows the user to enter whatever information they consider appropriate to highlight the main results. It is currently unstructured to allow maximum flexibility, bearing in mind that evidence synthesis can be based on the results table and the other structured information in the machine-readable record of the paper.

Secondary results table: This wizard allows the user to create a secondary results table to any specification, using measures that have been specified and any specified sub-samples.
Secondary results text: This wizard allows users to summarise or highlight results in the secondary results table. It is completely flexible and so also provide space for users to describe findings reported in additional results tables (see below).

Additional results table: This wizard allows any number of additional results tables to be created to describe additional findings. For example, this would be used to provide information about adverse events and process measures.

Research question answer: This wizard prompts users to report the answer to each of the research questions set out in the Introduction. This is used to populate the relevant section of the Discussion section. Failure to match answers to research questions with the research questions is a common failing of study reports; this system prevents that happening. It reminds the user of each question and prompts them to provide an answer that matches the question.

Theoretical implications: This wizard creates a field containing paragraphs in the Discussion section describing the implications of the study for the initially proposed logic model for the intervention (the mechanism by which the intervention was expected to have a larger effect than the comparator) and the theory behind it. Many RCT reports fail to make this kind of reference or do so in a way that is uninformative (e.g. saying 'the hypothesis was partially supported'). Building an incremental science of addiction involves much more than finding out whether interventions 'work' or not – we need a more specific articulation of theoretical predictions and the linkage between findings and these (https://psycnet.apa.org/record/2010-00152-001). This wizard provides a basis for doing that.

Practical implications: This wizard provides a free text field for users to convey the clinical or policy implications of the findings in the Discussion section, ensuring that these are not omitted from the report.

Limitations: This wizard populates a field in the Discussion section. It provides a set of prompts to ensure that users address all the types of limitation of the study, including sample representativeness, measurement error, generalizability from the sample, setting or intervention, choice of comparator etc.

Future research: This wizard populates a field in the Discussion section. It provides an opportunity for users to set out what they believe are research priorities for the future on the basis of the current findings.

Conclusions: This wizard populates the paragraph at the end of the Discussion section summarising the main conclusions of the study from the authors' perspective.

References: This wizard allows authors to enter references in free text but in due course will be linked to bibliographic software to be automatically generated by entries in other wizards.

Supplementary materials: This populates a section of the paper that describes the supplementary materials that will be submitted alongside the main manuscript.

Continued development of PAT

PAT is undergoing continual development to improve its functionality and scope in the light of users' experiences. Every wizard contains a 'Feedback' button to allow users to report issues or suggest improvements. This information is collated by the developers and discussed in order to prioritise updates.

Future variants of PAT will be able to be used for writing study protocols and even grant applications. In due course a PAT review tool should also be able to be used to pre-review articles to ensure that they conform to certain standards before submitting to a journal. This should make the process of submission and review much more efficient.

PAT and the Addiction Ontology

Underneath PAT is a preliminary version of an ontology covering the area of addiction (https://onlinelibrary.wiley.com/doi/10.1111/add.14554). This is a computable data structure that allows all the key entities that need to be mentioned in reports in the field to be classified and linked together. Ontologies are powerful tools for AI, allowing intelligent searching and inference and providing interoperability across diverse domains and communities of practice (https://link.springer.com/protocol/10.1007/978-1-4939-3743-1_1). The most successful example to date is the Gene Ontology (http://geneontology.org/) that has gone a long way to unifying the field of molecular biology. As the Addiction Ontology develops, it will provide more and more fine-grain intelligent prompting for information based on knowledge of the kind of study being reports and the topic of that study.