# How to write (and read) a paper

Sourendu Gupta*
*Department of Theoretical Physics,*
*Tata Institute of Fundamental Research,*
*Homi Bhabha Road, Mumbai 400005, India.*

A paper is a formal means of communication between scientists, whose purpose is to present a previously unknown result and a complete argument leading up to it. In this manual a top-down methodology is suggested as a tool to organize the writing or reading of a paper. As in any other top-down method, the first step is a short and precise statement of the result. Successive steps expand this until the paper is ready.

## I. THE NOTION OF A PAPER

A paper is a means of communication between scientists. It is meant to be readable for years: in some cases for over a hundred years [1]. In order to remain understandable for so long, a paper must have a formal structure, which changes little with fashion. Once you learn to use this structure you gain by being able to communicate with a large number of people over many years who are working on problems you are interested in. But, because the form is important, training is needed to read and write such a document.

All formal written scientific communications are not papers. One should distinguish full-length journal papers from their close relatives: letters, reviews and conference proceedings. This manual is about the writing of a full-length journal paper. You will need to write such papers so that others know what you have done. Your future career in science depends on the impact of your work, and papers are usually the medium through which such influences propagate. The means of assessing scientific impact are loose, but the number of citations is one of the ways which people use for this. Since a badly written paper is less likely to be read or used, and hence cited, you must learn how to write a paper well.

Modern papers evolved from informal communications between scientists. During the 16th and 17th centuries CE formal communication of results was through the medium of books. However, news of scientific advances was often communicated orally by travellers or in letters, usually reaching a large audience of peers before the books became available. Over the next centuries these letters became formalized instruments of communication. With the rise of scientific societies in Europe during the 18th century CE, records of oral presentations to peers became part of the formal process of communication of results. With the professionalization of science and the adoption of the system of tenure in universities, the notion of a scientific paper underwent modification, stabilizing by the end of the 19th century CE. The form of the paper (see Appendix A) has undergone little change since then, although the means of dissemination of papers has been revolutionized in the last few decades due to the rise of the internet.

In this manual I present a technique for writing a paper by successive expansion of a short core of the argument which is to be conveyed [2]. The steps in the construction of the paper can be reversed to deconstruct [3] and read other papers. The approach described here assumes that the decision to write a paper is taken after a clear result is obtained. Alternative approaches are sometimes espoused in other manuals [4]: the abstract is claimed to be discovered at the end of the paper, since one does not know what the paper might contain until it is written; such a paper is better off not being written. I do not discuss style, except very briefly in Appendix B.

## II. PLANNING THE PAPER

This section of the manual describes the process of planning the paper. Before you get to this point, you will have done a lot of work, discussed the material with co-workers, and come to a decision to write down your results as a paper. So, when you sit down to start writing, you have a reasonably complete idea of what you want to write. This manual will help you starting from this point.

In this section we will discuss how to construct three elements of the outline of the paper, namely a working abstract, a story-board and a working title. We devote one subsection to each of these elements. Once you construct this plan,

*Electronic address: `sgupta@tifr.res.in`

it would help to discuss this with your co-authors to check whether they agree with the detailed plan. There could be changes at this stage, but eventually all authors will agree on a common outline.

The notion of an outline of a paper is also very useful when you read one. The process of extraction of an outline from someone else's paper is so useful that I give it a name— deconstruction. It turns out that deconstructing a paper is a good way to understand it in a top-down fashion.

### A. The abstract: the bare bones of the paper

When you start to write a paper you probably have a very good idea of what you want to say. It is good to put this down in brief, as a reminder to yourself, and your co-authors. If you have the good fortune to have co-authors, sometimes you may be surprised at how many differences can arise over what your main result is.

If you write this core of the paper in a formal way then you have something which is known as an abstract. This comes at the beginning of a paper. It is a short and self-contained summary of the contents. It is normal practice for many scientists to scan through lists of abstracts to check whether they are sufficiently interested in a paper to read it. So the abstract is very important.

An abstract has a brief statement of (a) the problem being addressed, (b) the methods used and (c) the main results. It is a good idea to start the act of writing a paper by writing an abstract. This makes you think about what the primary and secondary results of your paper are, and how they fit into a subject. This is the briefest outline of your paper, and must be done first. You may come back to it later and polish it, but start with this.

The statement of the problem should be completed in one or two sentences, but still be sufficiently detailed that the reader will understand what you are working on. The only way to reconcile these two opposing requirements is to make assumptions about what the reader knows. Do not insult the reader by explaining the obvious, but do not assume so much that only three people in the world can follow your paper. The same remarks hold for the description of methods. It must be completed in about one sentence, and must make sense to the target audience. If you have made changes to well-known methods, then mention that.

The bulk of your abstract, which must be concise, should be a description of your results. A good rule of the thumb is to make this description as long as the other two parts of the abstract put together, *i.e.*, between two and four sentences. In order to do this, you really have to understand your work in its entirety, and distill the essence of it into the abstract. It is this process of thinking hard about your completed work which is important to the writing of the abstract, and a good paper.

If you have not analyzed the structure of a paper, then this is a good time to look hard at a few papers. No matter what the subject matter is, you will find that a modern paper begins with an abstract which has the three parts we described. Once you recognize this structure, it helps you to quickly understand what a paper is about. Here are a few examples. You can find and analyze more examples yourself.

#### *1. Example 1*

```
Serotonin (5-hydroxytryptamine or 5-HT) is thought to regulate
neurodevelopmental processes through maternalfetal interactions that
have long-term mental health implications. It is thought that beyond
fetal 5-HT neurons there are significant maternal contributions to fetal
5-HT during pregnancy, but this has not been tested empirically. To
examine putative central and peripheral sources of embryonic brain
5-HT, we used Pet1/ (also called Fev) mice in which most dorsal
raphe neurons lack 5-HT. We detected previously unknown differences in
accumulation of 5-HT between the forebrain and hindbrain during early and
late fetal stages, through an exogenous source of 5-HT which is not of
maternal origin. Using additional genetic strategies, a new technology for
studying placental biology ex vivo and direct manipulation of placental
neosynthesis, we investigated the nature of this exogenous source. We
uncovered a placental 5-HT synthetic pathway from a maternal tryptophan
precursor in both mice and humans. This study reveals a new, direct role
for placental metabolic pathways in modulating fetal brain development and
indicates that maternalplacentalfetal interactions could underlie
the pronounced impact of 5-HT on long-lasting mental health outcomes.
```

This abstract of a paper on developmental biology [5] is a classic example of a modern abstract (the lines are numbered for ease of analysis). The first two sentences, ending in line 5, set out the context and the problem. In lines 6–8 one method is discussed and the results from this are given in lines 8–11. In lines 11–13 a second method is described and the results from this are given in lines 13–15. Finally, the last sentence describes the overall conclusion of the paper.

### *2. Example 2*

`We present distance measurements to 71 high redshift type Ia supernovae discovered during the first year of the 5-year Supernova Legacy Survey (SNLS). These events were detected and their multi-color light-curves measured using the MegaPrime/MegaCam instrument at the Canada-France-Hawaii Telescope (CFHT), by repeatedly imaging four one-square degree fields in four bands. Follow-up spectroscopy was performed at the VLT, Gemini and Keck telescopes to confirm the nature of the supernovae and to measure their redshift. With this data set, we have built a Hubble diagram extending to z=1, with all distance measurements involving at least two bands. Systematic uncertainties are evaluated making use of the multi-band photometry obtained at CFHT. Cosmological fits to this first year SNLS Hubble diagram give the following results : Omega_M = 0.263 +/- 0.042(stat) +/- 0.032 (sys) for a flat LambdaCDM model; and w = -1.023 +/- 0.090(stat) +/- 0.054(sys) for a flat cosmology with constant equation of state w when combined with the constraint from the recent Sloan Digital Sky Survey measurement of baryon acoustic oscillations.`

The first sentence states the problem tackled in the paper. In lines 3–12 one finds a statement of the methods used. The rest of the abstract is essentially about the results. This abstract comes from an influential paper in astrophysics [6] which collected more than 1000 citations in 5 years.

### *3. Example 3*

`If high energy heavy ion collisions lead to the formation of a hot quark-gluon plasma, then colour screening prevents cc binding in the deconfined interior of the interaction region. To study this effect, the temperature dependence of the screening radius, as obtained from lattice QCD, is compared with the J/psi radius calculated in charmomium models. The feasibility to detect this effect clearly in the dilepton mass spectrum is examined. It is concluded that J/psi suppression in nuclear collisions should provide an unambiguous signature of quark-gluon plasma formation.`

This paper is another citation classic [7], with well over 1000 citations. It is also a classic example of a very clear abstract. The first sentence sets out the problem, the second the method, and the last two sentences give the results and conclusions.

### *4. Example 4*

`In a framework of the renormalizable theory of weak interaction, problems of CP-violation are studied. It is concluded that no realistic models of CP-violation exist in the quartet scheme without introducing any other new fields. Some possible models of CP-violation are also discussed.`

This is a remarkably elegant abstract from a paper which won its authors a Nobel prize [8]. Three brief sentences make up the three parts. The first sentence is the statement of the problem. The second and third sentences set out the conclusion. The part of the second sentence after the word ‘exist’ (see line 3) is the statement of the method! The method used is quantum field theory, but that phrase, and many others are never used. Such elegance and extreme compression is made possible by a superb understanding of the audience.

### B. The title: advertising your result

Once you have an abstract, you must compress it further into a title. The title should accomplish two things— first, it should say something about the scientific context, and, second, it should say something about the result. The title is like the slogan in an advertisement, so it pays to have an accurate and catchy title. Remember also that your paper will later be found through database searches, so you should try to include in the title one or two keywords which you think people might search by. You may revise your title at later stages, but do begin with a working title.

Here are examples of titles, good, bad, and ugly—

- “A Model of Leptons” [9] implies the context leptons and promises the description of a model. This is a superb advertisement for the solution of an outstanding problem which won a Nobel prize. Almost half a century later, we can only improve the title by changing “A” to “THE”.
- “Preparation of fatty acid methyl esters and dimethylacetals from lipids with boron fluoridemethanol” [10] has a detailed description of the context and implies that the novel element of the paper is a method of preparation. This is not a bad title, since it does its job correctly, but so clumsy that I find it a little ugly.
- “The Large N limit of superconformal field theories and supergravity” [11] just gives the context. It is a bad title for a work that has been one of the most influential in string theory.
- “Measurements of Omega and Lambda from 42 high redshift supernovae” [12] gives the context of high redshift supernovae with the result advertised as measurements of Omega and Lambda. The title works, and is both accurate and compressed enough that it is good.

### C. A story-board: organizing the argument

Now that you have decided which part of your research your paper will highlight, it is time to start putting the material together. Before you do any writing, it is good to organize your thoughts in some more detail. Film makers have the notion of a story-board which sets out the script pictorially; this allows them to easily move elements around until one has the best possible way of telling the story. This technique can be adapted to writing a paper— make your material take the place of pictures in a story-board.

This is easiest to do if you have lots of data. If you do, then elements of the story-board are figures and tables. During the course of your work you must have compiled some of these. Sit down with the complete set of figures and tables that you have, look at the abstract, and decide which ones illustrate your main conclusions. The story-board works very well even if your work is highly theoretical. In that case you probably have a lot of equations. Select the important ones for your story-board. Perhaps you have theorems; then the statements of the theorem are what you should have in your story-board. Sometimes it is useful to take small phrases, such as section titles, as part of the story-board. No matter what, think hard and long about what you would like to display on the story-board— these will be the highlights of your paper. Finally, when you have everything that you need, make sure that your notation is consistent and easy to follow.

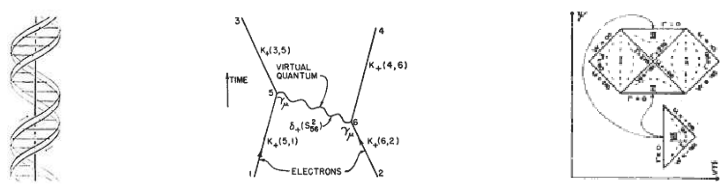


FIG. 1: Examples of pictorial representations of some highly influential ideas in science: (a) the double-helix structure of DNA [13], (b) certain integrals which arise in the treatment of the quantum theory of light and matter, now called Feynman diagrams [14], and (c) the causal structure of space-time, now called Penrose diagrams [15].

No matter how abstruse your result is, think hard about developing a pictorial representation of your primary conclusions. This is always helpful. Try to keep this final figure as nearly self-explanatory as possible, so that the main idea is clear even to people who do not read your paper carefully. Do not clutter the figure: make it convey the single most important idea. In the best of cases it will be imported by other people into their talks, become parts of reviews, and eventually become a cultural icon (like the some of the images in Figure 1). However, even if your figure does not reach the rarefied heights of cult status, time spent thinking of a good pictorial way to put across your main result is never wasted.

Once you have the elements of your story-board, shuffle them around to tell a clear and logical story. When you are satisfied, it is time to consult your co-authors with the working abstract, working title, and your story-board. Lively and interesting discussions usually ensue at this point, and often result in changes. At the end of this process of discussion you will have have an abstract, title, and story-board which all the authors agree on.

Film studios often have a first cut of a film which is shown to a small and selected audience. Their reactions are factored into the editing which results in a final cut. It is good to adapt this practise to your paper. Once you have a story-board which has the approval of your co-authors, try to give a seminar to your research group or department which uses the sequence you developed for the story-board. Feedback from your immediate colleagues can help you to rearrange the story-board in a more effective way. While preparing for a seminar you will find that the story-board *is* the seminar.

The rest of the job of writing will eventually reduce to connecting up the elements you put on the story-board. That is what we will take up in the next section.

### D. Reading a paper by deconstruction

The true test of a good tool is that it can be easily adapted to a slightly different purpose. The tool that we have built to write a paper certainly can be adapted immediately to construct a seminar based on your work. Less obviously, it also allows you to read a paper.

When you are confronted with a new paper, you already have a title and an abstract in front of you. In the previous sections we have seen how easy it is to parse these objects into their constituent parts, even if you know nothing of the subject. As a result, you know a little about the context of the paper, the method used and the main results obtained.

Now you want to go a little deeper and get an overview of the flow of the argument presented in the paper. Unlike the analysis of the title and abstract, this step requires understanding. You do this by deconstruction, *i.e.*, finding a story-board which made up the paper. Your story-board for the paper need not be the one that the authors used. However, if you read through the figures, tables and their captions, skim the formulae and their explanations, and read quickly through the conclusions, you will be able to construct a plausible story-board for the paper. This is your first reading of any new paper.

Here are some papers which you might want to deconstruct—

- If you are trying to read quantum field theory you will need to read [16], which won its author a Nobel prize. The paper is about the modern reformulation of quantum field theory, and is replete with figures and equations. Try to put the figures together with key equations in order to make a story-board for the paper.

- If you have any interest in polymers, you will read [17] some time. The paper has several figures and many equations. It also has a few appendices. Try to extract from the paper a story-board. Does the abstract help you in constructing a story-board?

- If your work involves numerical simulations, then you will read, sooner rather than later, the method which was first reported in [18]. The paper has a small number of tables, figures and equations. What story-board can you make out of these?

- If you are professionally interested in cancer, you will need to read this significant paper in the understanding of this pathology [19]. The first third of the paper is full of figures. The remainder has no figures, tables or equations. Extract a story-board from this paper and see which elements of the paper can contribute to a story-board.

Practice in deconstructing a paper is a good skill to develop. It helps you to skim the contents of a paper at a level somewhat deeper than reading the abstract. With practice, the time you will need to do such a deconstruction will become rather small. Once your have deconstructed a paper which you need to read, you will immediately see how different parts hang together. This will help you to plan how much time you need to spend in reading different parts of the paper. When your skill at deconstruction improves, you will find that preparing to give a talk in a journal club becomes a breeze. Eventually you will see that in your own area of expertise you will seldom need to do more, and when you do, you will quickly see the parts of the paper which you need to study more deeply.

## III. WRITING THE PAPER BY SUCCESSIVE EXPANSION

When you have a detailed outline you are ready to proceed further and complete writing the paper by successive expansion. Already, in going from the abstract to the story-board you used the method of expansion to put in details. Now we push this single method further.

### A. Captioning the story board

After you have the outline ready, the first step in writing the paper is to put down captions for the figures and tables. There are four kinds of information which may be given in a caption, of which only the first is mandatory—

- A typical caption must have a general description of the object.
- There may be an explanation of the quantities shown, and a key which explains things like colour coding, special notation, *etc.*.
- A caption may have a brief description of methods, or a reference to a more detailed explanation in the text or a citation.
- A caption may also point out some special feature of the data, or contrast it with some expectation (this part may have a reference to the text or a citation).

At the time you write the caption, the rest of the text is not yet written, so a reference to the text is a reminder to yourself that there is some explanation to be written down. Keep a list of such bits of the text: you will fill in these parts later. Here is where you also keep track of things like the notation which needs to be explained; put down such things in the list of bits described above. This is also the time at which you begin a separate bibliographic list, by putting in every citation that you need.

What about the story-board elements which are not figures or tables? Fill in brief details about these, like notation, special features, previous knowledge *etc.*. As you do this, build up the list of bits and the bibliography.

Finally, you may have put little phrases into your story-board. You could either write a small para expanding them a little, and in the process build up the list of bits and the bibliography, or you could add these items into the list of bits.

A word about citations: rest assured that many of your readers know the literature as well as you do, and some of them perhaps even better. Incomplete citations will lead them to believe that you are re-inventing the wheel. Cite everything which is important and relevant so that the reader knows that you know the state of the art, and that when you claim a new result, then it is very likely to be genuinely new.

The rest of the paper will be the expansion of the list of bits, the bibliography, and writing the logic which leads from one figure to another. Before proceeding to those tasks, it would be good to do some case analysis of captions.

There is, of course, a technique to captioning. This means that even if you are not very good at it to begin with, you can learn and become an expert. In order to become a real expert, examine and analyze as many captions as you can. To start you off I have collected a few examples and commented on them below.

#### *1. Example 1*

```
Surface and deep expression of a mesoscale eddy. (A) Inferred geostrophic
velocities at the surface and the associated time series of sea-level
anomalies (SLAs) from November 2004 to May 2005. Two eddies pass during
this time period: one in February to March during the period of bottom
observations and a second from the end of April to May, just after
the period of near-bottom observations. (B) Map of SLA on 21 February
2005 when an anticyclonic eddy made impact with the study site (white
star). The track of the center of the eddy from birth on 10 September
2004 until exit from the study region on 1 May 2005, marked every 4 weeks
(black circles), is superimposed on the map.
```

This is a caption to a two-part figure from [20]. The first sentence is a general description of both. The first sentence for each part is a general description of that part. The second sentence in each part points out a feature of interest. The phrases “mesoscale eddy” and “near-bottom observations” are clearly elements which will be explained in the text, and therefore would be in the list of bits.

### *2. Example 2*

```
Writhe distributions for lp = 50nm and lT = 74nm: PWLR(Wr) (green
dots), Pap(Wr) (red circles), and PWLC(Wr) (blue crosses). The chain
lengths are the following: (a) L = 2000nm, (b) L = 150nm, and (c) L
= 300nm. Note the different scales on the x-axis in the above figures.
```

This caption from [21] begins with a description of the content, and then gives a key to the figure. The second sentence gives supplementary information. The final sentence points out a feature of importance. The phrases “writhe”, “PWLR” and “Pap(Wr)” will clearly be part of the list of bits.

### *3. Example 3*

```
Inferred tree of phylogenetic relationships among mormyrid species
and morphs. The phylogeny was estimated by Bayesian analysis of cytb
sequences (values at nodes are posterior probabilities). A sequence from
the closest outgroup to the Mormyridae (Gymnarchus niloticus) was used
to root the tree. Green branches represent a small exterolateral nucleus
(EL) and magenta branches represent an enlarged EL divided into anterior
and posterior subdivisions (ELa/ELp); we reconstructed ancestral states
using parsimony (see text). Gray outline represents electric organs with
electrocyte stalks, and black outline represents electric organs with
developmentally labile stalks, based on a previous study (2).
```

The caption from [22] starts with a sentence of general description. The second sentence describes the method and the remainder is a detailed instruction of how to read the figure and is equivalent to a key). Note the embedded references to the text and citations. Apart from these, the phrases “Bayesian analysis”, “EL”, “ELa/ELp” will also be in the list of bits.

### *4. Example 4*

```
A self-reproducing automaton and a bacterium. (A) Von Neumanns logic. A
universal constructor (part A) constructs the hardware (part D) of the
offspring automaton according to the instruction I. The copier B makes
a copy of the software part of the offspring automaton, the instruction
tape I. Part A and part B are controlled by the regulatory part C. (B)
Phase contrast microscopy image of the division of an E. coli cell
(scale bar: 1 micro-m).
```

This caption extracted from [23] is another two-part figure. The first sentence is a general description of the two together. The first sentence of each part is a general description of that part. The remainder is a detailed explanation of the figure. The phrases “automaton”, “universal constructor” and “copier” will be on the list of bits.

### *5. Example 5*

```
In this calculation it is unnecessary to choose an explicit set of
coordinates q for the classical phase space. (Such a choice would be
very awkward in the nonlinear sigma model because of the nonlinearity
of the phase space.) It is enough to have a basis of tangent vectors
```

to the phase space (analogous to the tetrad in general relativity). The matrices F-tilde and F may be constructed relative to any such basis. In the non-linear sigma model a very convenient basis of tangents to the phase space are the matrices g-inverse times delta g. The matrix F must act both on the Lie algebra index of g-inverse delta g and on sigma.

This paragraph from [24], which deals with a mathematical treatment of certain quantum field theories, serves the same purpose as a caption; it explains a certain computation which is displayed after this. The quantities $\tilde{F}$, $F$, $\sigma$ and $g^{-1}\delta g(\sigma)$ are clearly on the list of bits to be explained.

### B. The definitions and methods sections

These two parts are now straightforward. The first attempt to write these sections will start from the list of bits you have put together while writing the captions. Put together little explanations of each of these items. When your explanation is done, take the item off the list. This explanation could be a single sentence like "X was earlier defined in [citation C]" or a short paragraph of technical explanation. Unless you are developing a new subject from the ground up, most of the definitions are fairly standard, and you just need to cite the correct papers to point people at them. There may be just a few definitions which you need to extend or re-interpret in some way. Put these down carefully; add in the reasons for making these redefinitions. If you need to add new terminology to explain some of this stuff, then that goes into the list of bits remaining. When the list of bits is empty, you have explained everything that you need to.

This is the core of the definitions and methods sections. Go back and take a look at your story-board. Some of the elements there (equations, figures, tables) may be about methods. Move these into the methods section. Read through what you have written once more, polish it up if necessary, change the order to a more logical one, polish it up some more, and, before you know it you are done with a first draft of the definitions and methods sections!

### C. The results section

One of the best pieces of advice about writing papers that I received from my doctoral adviser was to leave the beginning to the end [25]. Do that. You already have the definitions and methods sections. If you leave out the introduction and the conclusions, then the rest of the paper is a set of observations leading from one result to another. Describe these in sequence. If a step in the argument needs a result established elsewhere, then add a citation to that paper.

The results section of your papers will build the core of your scientific reputation. The logical presentation of solid results, is what you should aim for, while pointing out what kind of care was taken to control unknowns and how loopholes in logic were plugged, You should also take on the role of a guide to your readers, as you take them on a tour of your story-board. You must explain why you stop at a certain element of the story, and then discuss that element in sufficient detail. Then you should point out where you will take them next, and why.

One of the most embarrassing things that can happen to you is that years after your wrote a paper, and when your interests have changed completely, someone comes up to you and asks you to explain one part of it. You could have forgotten that detail entirely, or you may need quite a bit of time to reconstruct that argument. To forestall such events, write down every step: concisely, of course, but without omission.

Interesting observations which are immaterial to the main development of your paper must go into footnotes or parenthetical paragraphs. Possibly the most famous such remark is the following sentence from the paper on the structure of DNA, "It has not escaped our notice that the specific pairing we have postulated immediately suggests a possible copying mechanism for the genetic material" [13]. Material of this sort should not interrupt the main flow of your argument.

### D. The conclusions

When I was a young post-doc I was trying to convince a famous physicist that work which he had attributed to someone else was actually done earlier in one of my papers. He told me something which I've come to realize is true of many others, "When I read a paper, I go directly from the abstract to the conclusions". In doing this people may miss the exciting details, and the real nitty-gritty of research, nevertheless this is what some of them do. The reason is that the formal strcture of a conclusion allows them to take a dangerous short-cut to understanding.

The conclusion will contain a quick account of the logic of your arguments and can refer back to the main storyboard elements (tables, figures, equations and section headings). Some speculation is allowed here. Do not worry about being too outrageous. If you are, then the referee will surely point it out, unless your co-authors beat you to it.

Here is a typical modern concluding section from the paper which reported the first experimental realization of quantum teleportation (dubbed the "Beam me up, Scotty" experiment by newspapers) [26]. Note how the first two paragraphs mix an account of the results reported in the paper with suggestions for new directions of work. The last paragraph is entirely speculative.

> In our experiment, we used pairs of polarization entangled photons as produced by pulsed down-conversion and two-photon interferometric methods to transfer the polarization state of one photon onto another one. But teleportation is by no means restricted to this system. In addition to pairs of entangled photons or entangled atoms, one could imagine entangling photons with atoms, or phonons with ions, and so on. Then teleportation would allow us to transfer the state of, for example, fast-decohering, short-lived particles, onto some more stable systems. This opens the possibility of quantum memories, where the information of incoming photons is stored on trapped ions, carefully shielded from the environment.
>
> Furthermore, by using entanglement purification---a scheme of improving the quality of entanglement if it was degraded by decoherence during storage or transmission of the particles over noisy channels---it becomes possible to teleport the quantum state of a particle to some place, even if the available quantum channels are of very poor quality and thus sending the particle itself would very probably destroy the fragile quantum state. The feasibility of preserving quantum states in a hostile environment will have great advantages in the realm of quantum computation. The teleportation scheme could also be used to provide links between quantum computers.
>
> Quantum teleportation is not only an important ingredient in quantum information tasks; it also allows new types of experiments and investigations of the foundations of quantum mechanics. As any arbitrary state can be teleported, so can the fully undetermined state of a particle which is member of an entangled pair. Doing so, one transfers the entanglement between particles. This allows us not only to chain the transmission of quantum states over distances, where decoherence would have already destroyed the state completely, but it also enables us to perform a test of Bell's theorem on particles which do not share any common past, a new step in the investigation of the features of quantum mechanics. Last but not least, the discussion about the local realistic character of nature could be settled firmly if one used features of the experiment presented here to generate entanglement between more than two spatially separated particles.

### E. The introduction

You are now ready to tackle the most challenging part of the paper— the introduction. This is the place where you point out the context of the paper and showcase your scholarship. There is a certain logic to a subject, and any new work rearranges that logic to some extent. The introduction is the only place where this very wide view of a subject appears. While well-written results and conclusions sections build your reputation amongst peers, well-written introductions influence newcomers to the field, and thereby seal your reputation. Remember Newton's remark about standing on the shoulders of giants. The introduction lets other people know whether you are comfortable on such a precarious perch, or tend to slip and fall. So give this your best shot, but only when you are done with the rest of the paper.

Here is an introduction I especially treasure [27]. In a four paragraph communication, not quite a paper, and shorter even than most letters, the authors first set out the context of the work brilliantly. At this extreme level of compression, the distinction between an abstract and an introduction blurs, but this paragraph does a wonderful job of both.

> In a paper under this title which has recently appeared, we have described and discussed observations which have led us to the conclusion that the light quantum possesses an intrinsic spin equal to one Bohr unit of angular momentum. In the four weeks which have elapsed since that paper was put into print, the experimental technique has been much improved in the direction of attaining greater precision. It appears desirable forthwith to report our newer results, which confirm the conclusion stated above.

## IV. SUMMARY

There are many ways to approach the writing of a paper. Here I have set out a top-down method of writing a paper by successive expansion of a core argument. The process starts by putting down a very compressed account of the paper in the form of an abstract and a title. This is followed by an expansion of the argument into a story-board, which consists of the main results to be displayed in the paper. These elements together comprise the plan of the paper. Such a plan is easily adapted to the job of creating a seminar presentation, and is almost exactly the presentation itself.

The plan is then expanded into the paper using the same top-down method by filling in captions, links in the arguments which lead between the elements of the story-board and so on, until all the steps in the argument are either written down in detail or delegated to cited papers. This leaves the writing of the conclusions and the introduction, about which I say a few words.

The inverse process of reading a paper is one I have named deconstruction. It consists of orienting yourself by the abstract and title and using these to extract a story-board from a paper. Just as a story-board of your own paper needs little expansion to make a seminar, the story-board of someone else's paper is also almost enough material to make a seminar out of. Once a story-board is extracted for a paper being read, it becomes clear which parts of the argument are general background knowledge in the field, and which are specific advances due to the paper under study. Practice in deconstruction of papers is an essential skill for any scientist.

To repeat, a top-down methodology is very well suited for writing papers. Various parts of papers have very formal structures, which have been described in this manual. A central process in the reading and writing of papers is the creation of a story-board. This is hidden structure in a paper, which is elaborated while writing, and needs to be discovered when reading.

### Appendix A: The parts of a paper

A full-length research paper consists of the following elements, listed in the order in which they appear—

1. The title, which is a short description of the subject of the paper.

2. A list of authors and their affiliations. In this manual I do not discuss this aspect of the paper.

3. The abstract, which is a short outline of the argument of the paper. It contains a brief description of the context of the work, the methods used and the major conclusions reached.

4. Classification keys used for electronic information retrieval, such as report Numbers, keywords or formal index items, for example, PACS (Physics and Astronomy Classification Scheme) numbers. These are becoming less important as algorithms for extracting these elements directly from the paper improve.

5. The body of the paper, which is divided into several sections. Usually

    (a) The introductory section, which is a short survey of the field, whose purpose is to outline the context in which the work is being reported.

    (b) The definitions and methods sections, which present detailed definitions and describe clearly, completely, but concisely, the methods used in the work. The two sections may be merged into one if either is short enough.

(c) The results section, which should normally display and state the results with objectivity and the minimum of interpretation (a task which is delegated to the conclusions).

(d) The conclusions, which gives a brief outline of the conclusions reached, places them into context of previous knowledge, and may speculate on new directions opened out by the work.

6. Acknowledgements, if any. It is often necessary to acknowledge grants and funding bodies. There is no discussion of this aspect of papers in this manual.

7. The list of citations, cross linked into the body of the paper.

**Appendix B: Style**

The style of a paper is eventually your own, but most journals impose some constraints. While some of these constraints vary from one journal to another, here are some points which are very common—

1. The language used should be simple in order to reach far in space and time. English is the common language of science in our era, and therefore used by scientists with different levels of skill in the language.

2. Discussions are not descriptions. When you present the discussion of a result, it should not merely be a description of a graph which is presented. Ruthlessly remove such sentences. Consider what more you have to say.

3. Avoid the subjective. If you are unsuccessful with an approach, simply say that, instead of “We felt method A is unsuitable”. If you have an argument that method A is indeed unsuitable, then give it. But do not express your feelings about the method.

4. Delete adjectives. After giving vent to your elation by writing “Exciting conclusions about X were reached”, quietly go back and delete the word “exciting”. The most prestigious journals will actually do this for you if you don’t have the heart to.

5. The passive voice was once widespread, but its dominance is eroding. “We observed that $\cdots$” is more modern than “It was observed that $\cdots$”. If you are uncomfortable with the singular “I observed that $\cdots$”, do not erroneously substitute “we” for “I”. Instead, consider moving to the passive voice.

Style manuals of many journal are publicly available [28].

---

[1] S. Arrhenius, *Phil. Mag. and J. Sci.*, Ser 5, 41 (1896) 237, has recently become a widely read reference in climate science.
[2] This is an extrapolation of a method developed to communicate with computers in N. Wirth, *Comm. ACM*, 14 (1971) 221.
[3] While this word and the process it describes may have some superficial similarity to the sense in which it is used in literary criticism, the aim here is different. For a partial explanation of the meaning that this term has in criticism, see *Letter to a Japanese Friend* by J. Derrida in *Derrida and différance*, by D. C. Wood and R. Bernasconi, Northwestern University Press, 1988 (also available at http://lucy.ukc.ac.uk/Simulate/Derrida_deconstruction.html)
[4] See, for example, the Rice biolabs manual: http://www.ruf.rice.edu/ bioslabs/tools/report/reportform.html
[5] A. Bonnin *et al.*, *Nature*, 472 (2011) 347.
[6] P. Astier *et al.*, eprint astro-ph/0510447.
[7] T. Matsui and H. Satz, *Phys. Lett.* , B 178 (1987) 41.
[8] M. Kobayashi and T. Maskawa, *Prog. Theor. Phys.* , 49 (1973) 652.
[9] S. Weinberg, *Phys. Rev. Lett.* 19 (1967) 1264.
[10] W. R. Morrison and L. M. Smith, *J. Lipid Res.*, 5 (1964) 600.
[11] J. Maldacena, *Adv. Theor. Math. Phys.*, 2 (1998) 231.
[12] S. Perlmutter *et al.*, *Astrophys. J.* 517 (1999) 565.
[13] J. D. Watson and Crick, *Nature*, 171 (1953) 737.
[14] R. P. Feynman, *Phys. Rev.* 76 (1949) 769.
[15] B. Carter, *Phys. Rev.* 141 (1966) 1242.
[16] K. G. Wilson, *Phys. Rev.* , D 10 (1974) 2445.
[17] P. G. de Gennes, *J. Chem. Phys.*, 55 (1971) 572.
[18] N. Metropolis *et al.*, *J. Chem. Phys.*, 21 (1953) 1087.

[19] J. F. R. Kerr, C. M. Winterford, B. V. Harmon, *Cancer*, 73 (2006) 2013.
[20] D. K. Adams *et al.*, *Science*, 332 (2011) 580.
[21] S. Medalion, S. M. Rappaport and Y. Rabin, *J. Chem. Phys.* 132 (2010) 045101.
[22] B. A. Carlson *et al.*, *Science*, 332 (2011) 583.
[23] V. Noireaux, Y. T. Maeda, A. Libchaber, *Proc. Nat. Acad. Sci.*, 108 (2011) 3473.
[24] E. Witten, *Comm. Math. Phys.* 92 (1984) 455.
[25] B. Banerjee, private communication.
[26] D. Bouwmeester *et al.*, *Nature*, 390 (1997) 575.
[27] C. V. Raman and S. Bhagavantam, *Nature*, 129 (1932) 22.
[28] http://www.nature.com/nphys/journal/v3/n9/full/nphys724.html
https://authors.aps.org/STYLE/