Thinking about Abstracts

John Carlis

Computer Science and Engineering, CS&E
Biomedical Informatics and Computational Biology, BICB
University of Minnesota
carlis@umn.edu

May, 2014

Abstract

Every research paper's abstract briefly tells its story. Unfortunately, many writers produce a poor abstract, often in part looking like a table of contents, thereby lessening the chance of getting the research published and weakening its eventual impact too. This missive will help you improve your abstract reading and writing skills. First, you will reverse engineer an existing abstract, succinctly stating each abstract sentence's purpose. By repeating this exercise you will see storyline patterns, and be able to determine why you like or dislike an abstract. Second, you will polish an abstract so if fits into the publishing venue's word and perhaps sentence limits.

While improving your conception of abstracts is useful, your ultimate goal is to readily write an abstract mimicking the style and storyline of other abstracts written for your audience, and employing a sensible level of detail, thereby gaining that audience's approval. As with learning any skill, you must thoughtfully practice, but it is worth the effort.

1. Introduction[1]

Suppose you wanted to write a sonata. What would you do? Well, of course, you would listen to sonatas, not just for pleasure, but also to find out what makes a sonata a sonata. You would look for what to include and exclude. You would look for elements fitting or not fitting sonata patterns. You would, in learning theory terms, acquire a mental map of the sonata territory and be able to move around in it.[2] You would get past being, sometimes, a lucky writer, learning to see patterns and then deliberately controlling them.[3]

You also would study what others have written about writing sonatas, but, of course, would not expect to learn too much. Why not? Well, no book ever teaches any real skill – you have to work at it. Indeed, learning any skill combines: following instructions, doing graduated exercises, and getting feedback (or muddling through on your own). Improving includes slowing down, breaking a whole skill into its parts, practicing basics until they become second nature, practicing whole skills in controlled contexts, just doing it, talking about it later, listening to elder wisdom [4] [the meandering babble of washed up used-to-be's ☺] – taking it all in, practicing, getting better and practicing, followed by more practicing, because craft requires practice.

What's an abstract? It is a brief, self-contained statement conveying to the reader the paper's essence, so she/he can decide whether to read the whole paper. But: what should your abstract say and not say? Here are several general thoughts.

First, read a bunch of published abstracts. Then mimic them – in style. Indeed, if your style is familiar to readers then they will have an easier time understanding your substance. Of course, mimicking is not slavery. You do not have to craft a dull abstract just because others have.

---

[1] So, since a reader has just read the abstract, how should I begin? Hmmm, I'll try an analogy about learning which will make me appear cultured and not geeky. I bet Helen Sword, author of "Stylist Academic Writing" (Harvard University Press, 2012), would approve.
[2] Read G. Caine and R. Caine, "Making Connections: Teaching and the Human Brain," Dale Seymour Publ., 1994.
[3] When you can "see", you can draw. Read Betty Edwards, "Drawing on the Right Hand Side of the Brain" Tarcher, 1999. It's a great book. The larger point here is: construct a mental model of abstracts.
[4] Read Gary Klein's "Sources of Power: How People Make Decisions", MIT Press, 1999.

Second, a helpful place to look is the "instructions for authors" that just about any journal, and many conferences, make available. Find those instructions before you begin to write, and follow them.

Third, "read like a writer"[5], so that you perceive good and bad patterns in writing, and then, respectively, use and avoid them in your writing. I think of "read like a writer" in two ways. I think of it as reverse engineering, where I take apart finished abstracts and study them. I look for common elements. I make up a vocabulary for naming elements, and then use that vocabulary as I construct my own. I also think of it as like unconsciously absorbing a culture, trying to fit in, and gradually becoming less of an alien. For example, think of somebody you know who has English as a second language, or started playing basketball as an adult. Usually, you can easily distinguish them from native speakers or gym rats. Some people, however, mimic well, and quickly fit in. When you read like a writer you find things to mimic and you appear to be a genuine scholar. When you read like a writer look for (and not mimic) a flaw that many abstract exhibit, namely, a table of content flavor. It is a flaw because an abstract should be self-contained, and something like "We discuss two issues …" cannot be understood without reading the body of the text, and so is not self-contained.

The rest of this missive gives more specific advice. The next two sections, "Purpose" and "Polish" examine one pretty good research paper abstract and then modify it. Purpose and polish are skills you need for writing abstracts. Sorry, there is no universal storyline for an abstract. Nor is there one way to polish. You need to work. Fortunately, you are quite bright and surely you can learn this skill. Section 4 follows with several more abstracts. Section 5 summarizes the abstract storylines (the sentence purposes) generated by me and by graduate students attending workshops. Those several dozen purpose and polish examples are in other, "example annex" documents. Feel free to study them.

Because I teach workshops on writing to them, this missive targets graduate students [plus postdocs and faculty] in Computer Science, CS, and Biomedical Informatics and Computational Biology, BICB. However, the advice here is, I think, broadly applicable.

Warning! As with any skill you can't be passive, you must work. [You do not get fit merely watching an exercise video!] So, when you see "Work" in the following sections – do it.

2. Purpose.

Here an essential point: a good abstract tells a story. Furthermore, the parts of the story appear in a sensible order, one commonly seen by its (limited) audience. It, therefore, meets its readers' expectations.[6] When writing my own abstracts or helping others, I find that stating the purpose(s) of each sentence helps me see the story's flow from a reader's point of view, and readily find flaws in the story. If you think of an abstract as telling a story, then of course only some content and orderings are effective, and you can judge when content and order works or doesn't.[7] To help me study I put each sentence on a new line. (I recommend this tactic of spreading out any paragraph's text in my One-Draft manuscript; see page 13).

To begin, study the abstract below, which is pretty good – the paper was published. Here I formatted the abstract into a table, and added a blank column. Notice that you have some work to do. My phrases follow on the next page, but don't cheat and look before writing your own.

WORK: fill in the left column with each sentence's purpose. Use one, or just a few words for each purpose – be succinct. You should expect to struggle a bit in finding good purpose phrases, after all it's a new task.

Warning! Don't think that if you do not know anything about the substance of this abstract that you cannot perform the work. You can; try it.

"A study of hierarchical and flat classification of proteins", Zimek A, Buchwald F, Frank E, Kramer S., IEEE/ACM Trans Comput Biol Bioinform. 2010 Jul-Sep;7(3):563-71.

---

[5] Google "read like a writer". You'll find many sites.
[6] See the George Gopen reference in my One Draft manuscript [www.cs.umn.edu/~carlis]
[7] Read Robert Pirzig's "Zen and the Art of Motorcycle Maintenance", Harper Collins, 1974. Look for the passage about judging essays.

Abstract

	Automatic classification of proteins using machine learning is an important problem that has received significant attention in the literature.
	One feature of this problem is that expert-defined hierarchies of protein classes exist and can potentially be exploited to improve classification performance.
	In this article, we investigate empirically whether this is the case for two such hierarchies.
	We compare multiclass classification techniques that exploit the information in those class hierarchies and those that do not, using logistic regression, decision trees, bagged decision trees, and support vector machines as the underlying base learners.
	In particular, we compare hierarchical and flat variants of ensembles of nested dichotomies.
	The latter have been shown to deliver strong classification performance in multiclass settings.
	We present experimental results for synthetic, fold recognition, enzyme classification, and remote homology detection data.
	Our results show that exploiting the class hierarchy improves performance on the synthetic data but not in the case of the protein classification problems.
	Based on this, we recommend that strong flat multiclass methods be used as a baseline to establish the benefit of exploiting class hierarchies in this area.

Here are my sentence purposes. Do you agree with them?

context	Automatic classification of proteins using machine learning is an important problem that has received significant attention in the literature.
opportunity[8]	One feature of this problem is that expert-defined hierarchies of protein classes exist and can potentially be exploited to improve classification performance.
approach	In this article, we investigate empirically whether this is the case for two such hierarchies.
method	We compare multiclass classification techniques that exploit the information in those class hierarchies and those that do not, using logistic regression, decision trees, bagged decision trees, and support vector machines as the underlying base learners.
crux	In particular, we compare hierarchical and flat variants of ensembles of nested dichotomies.
why use	The latter have been shown to deliver strong classification performance in multiclass settings.
kind of evidence	We present experimental results for synthetic, fold recognition, enzyme classification, and remote homology detection data.
results	Our results show that exploiting the class hierarchy improves performance on the synthetic data but not in the case of the protein classification problems.
impact	Based on this, we recommend that strong flat multiclass methods be used as a baseline to establish the benefit of exploiting class hierarchies in this area.

What purpose terms?

There is no hard and fast set of purpose terms, and any pair of abstracts need not follow the same storyline. Fortunately, after you reverse engineer several dozen abstracts you will settle on terms that you will use most of the time, and you will discern common storylines. (See Section 5.)

WORK

To improve you need to practice. Here are your work tasks:

- **Pick a journal that you read and might publish in.**
- **Reverse engineer a bunch of abstracts.**
- **Grow a list of purposes [it will evolve as you work].**
- **Look for common storylines, i.e., sequences of purpose.**
- **When you find an abstract that is an outlier or flawed in some way rewrite it to fit a common pattern.**
- **Pick a journal from a field far from yours and see how similar and different their abstracts are.**

[8] After the context readers generally expect to see a **problem** or an **opportunity**.

3. Polishing.

Here is a situation I hope you face often in your career. Suppose you send a paper to a top-flight journal, say Nature Methods.[9] An abstract for a Nature Methods paper varies in length depending on the kind of paper. A full research article has a 150 word limit, while a brief communication has a 70 word limit. If you submitted the former and the editors said they like it and will publish it, but only as the latter, then you'll have polishing work to do. You could object, claiming "every one of my words is a priceless pearl, and is inviolate." You could reluctantly accede and then discard whole sentences, and produce a poor story. Better, of course, is to say “sure” and do some wordsmithing[10].

3.1. Initial polishing WORK: With just a bit of wordsmithing, reduce the abstract from by about ~10%. Then explain what wordsmithing you did. By naming these editing patterns you acquire wordsmithing skills. DO NOT read my edits before doing your own.

purpose	original [181 words]	your revision
context	Automatic classification of proteins using machine learning is an important problem that has received significant attention in the literature.	
oppor-tunity	One feature of this problem is that expert-defined hierarchies of protein classes exist and can potentially be exploited to improve classification performance.	
approach	In this article, we investigate empirically whether this is the case for two such hierarchies.	
method	We compare multiclass classification techniques that exploit the information in those class hierarchies and those that do not, using logistic regression, decision trees, bagged decision trees, and support vector machines as the underlying base learners.	
crux	In particular, we compare hierarchical and flat variants of ensembles of nested dichotomies.	
why use	The latter have been shown to deliver strong classification performance in multiclass settings.	
kind of evidence	We present experimental results for synthetic, fold recognition, enzyme classification, and remote homology detection data.	
results	Our results show that exploiting the class hierarchy improves performance on the synthetic data but not in the case of the protein classification problems.	
impact	Based on this, we recommend that strong flat multiclass methods be used as a baseline to establish the benefit of exploiting class hierarchies in this area.	

Not too tough is it?

[9] See http://www.nature.com/nmeth/pdf/gta.pdf

[10] A wordsmith is a skillful writer – that's you. For example, replace "do some wordsmithing" with simply "wordsmith."

See what you think of my revisions. [Since I might have misinterpreted the authors' intension, I checked with them and the revised one is ok.] Notice that here I edited only **within** a sentence. Later you will see examples of **combining and reordering** sentences.

purpose	original [182 words]	my revision [161 words]
context -8	Automatic classification of proteins using machine learning is an important problem that has received significant attention in the literature.	Classifying proteins is an important problem in bioinformatics and computational biology.
opportunity -4	One feature of this problem is that expert-defined hierarchies of protein classes exist and can potentially be exploited to improve classification performance.	Since expert-defined hierarchies of protein classes exist, machine learning techniques potentially can exploit them to automatically classify proteins.
approach -2	In this article, we investigate empirically whether this is the case for two such hierarchies.	Here, we investigate empirically whether this is the case for two such hierarchies.
method no change	We compare multiclass classification techniques that exploit the information in those class hierarchies and those that do not, using logistic regression, decision trees, bagged decision trees, and support vector machines as the underlying base learners.	We compare multiclass classification techniques that exploit the information in those class hierarchies and those that do not, using logistic regression, decision trees, bagged decision trees, and support vector machines as the underlying base learners.
crux no change	In particular, we compare hierarchical and flat variants of ensembles of nested dichotomies.	In particular, we compare hierarchical and flat variants of ensembles of nested dichotomies.
why use -1	The latter have been shown to deliver strong classification performance in multiclass settings.	The latter have delivered strong classification performance in multiclass settings.
kind of evidence no change	We present experimental results for synthetic, fold recognition, enzyme classification, and remote homology detection data.	We present experimental results for synthetic, fold recognition, enzyme classification, and remote homology detection data.
results -2	Our results show that exploiting the class hierarchy improves performance on the synthetic data but not in the case of the protein classification problems.	Our results show that exploiting the class hierarchy improves performance on the synthetic data but not in the case of classifying proteins.
impact -2	Based on this, we recommend that strong flat multiclass methods be used as a baseline to establish the benefit of exploiting class hierarchies in this area.	Therefore, we recommend that strong flat multiclass methods be used as a baseline to establish the benefit of exploiting class hierarchies to classify proteins.

Notice that I prefer "we", first person. and active voice. Check to see what your journal says.[11]

[11] Helen Sword, In "Stylish Academic Writing" (Harvard Univ. Press, 2012) points out that many journals not encourage first person and active voice, but some reviewers and writers baulk.

3.2 More polishing

For now let's suppose you were lucky and the journal imposed a 150 word limit. **Work**: remove about another 11 words. Think about what is most important to keep.

purpose	**my first revision [161 words]**	**your revision**
context	Classifying proteins is an important problem in bioinformatics and computational biology.	
opportunity	Since expert-defined hierarchies of protein classes exist, machine learning techniques potentially can exploit them to automatically classify proteins.	
approach	Here, we investigate empirically whether this is the case for two such hierarchies.	
method	We compare multiclass classification techniques that exploit the information in those class hierarchies and those that do not, using logistic regression, decision trees, bagged decision trees, and support vector machines as the underlying base learners.	
crux	In particular, we compare hierarchical and flat variants of ensembles of nested dichotomies.	
why use	The latter have delivered strong classification performance in multiclass settings.	
kind of evidence	We present experimental results for synthetic, fold recognition, enzyme classification, and remote homology detection data.	
results	Our results show that exploiting the class hierarchy improves performance on the synthetic data but not in the case of classifying proteins.	
impact	Therefore, we recommend that strong flat multiclass methods be used as a baseline to establish the benefit of exploiting class hierarchies to classify proteins.	

Polishing is getting harder, but still not so bad?

3.3 Oh, No! Even more polishing

Now suppose you were not so lucky and the journal imposed a <u>100 word limit</u>.

Work: meet that new word limit. Now you will really have to think about what is most important to keep. You may end up revising the purposes too.

My version follows – don't cheat; do yours first.

purpose	your previous revision	your new revision
context		
opportunity		
approach		
method		
crux		
why use		
kind of evidence		
results		
impact		

My version [92 words]

13 context; opportunity	Machine learning techniques potentially can exploit expert-defined protein hierarchies to classify proteins automatically.
30 approach; method; things we did	Here, employing logistic regression, decision trees, bagged decision trees, and support vector machines as underlying base learners, we compare previously-successful flat variants against two different hierarchical ensembles of nested dichotomies.
28 results	We find that, for synthetic and protein (fold recognition, enzyme classification, and remote homology detection) data, exploiting the hierarchy improves performance for synthetic but, unfortunately, not protein data.
21 impact	Therefore, we recommend employing strong flat multiclass techniques as a baseline when establishing the benefit of exploiting hierarchies to classify proteins.

Writing Sidebar

Gratuitous synonyms cause problems in technical writing, making a passage harder to understand and longer than it needs to be. Pick **one** term for a notion and use only it! Here I discovered that "technique" and "method" were synonyms. Learn to look for such things. When you choose one term then you will more readily see redundancies that you can squeeze out. You will also confuse your readers less often.

3.4 Oh, No again! Even more polishing

Now suppose you were not so lucky and the journal imposed a 70 word limit. **Work**: meet that new word limit. Now you will really have to think even harder about what is most important to keep. You may end up revising the purposes too.

purpose	your previous revision	your new revision

My even shorter version [72 words]:

13 context; opportunity	Machine learning techniques potentially can exploit expert-defined protein hierarchies to classify proteins automatically.
30 -12 method	Here, employing underlying base learners, we compare previously-successful flat variants against two different hierarchical ensembles of nested dichotomies.
28 results	We find that, for synthetic and protein data, exploiting the hierarchy improves performance for synthetic but, unfortunately, not protein data.
21 -8 impact	Therefore, we recommend employing strong flat multiclass techniques as a baseline when establishing the benefit of exploiting hierarchies to classify proteins.

Notice that I removed these detail words:
logistic regression, decision trees, bagged decision trees, and support vector machines, fold recognition, enzyme classification, and remote homology detection)

Work: Remove 2 more words.

Work: Is the abstract still ok or did I lose the story?

3.5 Oh, No yet again! Even more polishing

Ooops, I forgot that the 70 word limit also included a 3 sentence limit. **Work**: meet that new word/sentence limit.

Again: don't cheat and look ahead.

purpose	your previous revision	your new revision

My even shorter version [3 sentences; 69 words]:

purpose	my revision
28 context; opportunity method	Since machine learning techniques, exploiting expert-defined hierarchies, might successfully classify proteins automatically, we compared, employing underlying base learners, previously-successful flat variants against different hierarchical ensembles of nested dichotomies.
20 results	We find that, for synthetic and protein data, exploiting the hierarchy improves performance for synthetic but, unfortunately, not protein data.
21 impact	Therefore, we recommend employing strong flat multiclass techniques as a baseline when establishing the benefit of exploiting hierarchies to classify proteins.

Finally, Kramer, the senior author, gave me a succinct summary, which is too terse for an abstract but works as a talk's conclusion. He said "Don't use hierarchies".

Purpose and Polish Summary

You need to analyze lots of abstracts, that is, thoughtfully practice! So, perform this sentence purpose and polishing on dozens of abstracts. You will learn a lot if you can articulate the flaws in a poor abstract and then improve it. Get and study writing books too. Here is my favorite.

Gerald Alred, Charles Brusaw & Walter Oliu, "Handbook of Technical Writing," 9th Edition, St. Martin's Press, 2008. *Hint: look at "conciseness"*

4. More examples.

This Section contains several more examples. **Do not cheat and look at my version first.**

WORK: Take this published, pretty good abstract. First determine each sentence's purpose, and then reduce the abstract by a third, that is, to **100 words** or fewer.

"Meta Optimization and its Application to Portfolio Selection", Puja Das and Arindam Banerjee, KDD'11, August 21–24, 2011, San Diego, California, USA

purpose	original [151 words]
	Several data mining algorithms use iterative optimization methods for learning predictive models.
	It is not easy to determine upfront which optimization method will perform best or converge fast for such tasks.
	In this paper, we analyze Meta Algorithms (MAs) which work by adaptively combining iterates from a pool of base optimization algorithms.
	We show that the performance of MAs are competitive with the best convex combination of the iterates from the base algorithms for online as well as batch convex optimization problems.
	We illustrate the effectiveness of MAs on the problem of portfolio selection in the stock market and use several existing ideas for portfolio selection as base algorithms.
	Using daily S&P500 data for the past 21 years and a benchmark NYSE dataset, we show that MAs outperform existing portfolio selection algorithms with provable guarantees by several orders of magnitude, and match the performance of the best heuristics in the pool.

WORK: Take this published, pretty good abstract. First determine each sentence's purpose, and then reduce the abstract to **50 words** or fewer.

"Bridging Taxonomic and Disciplinary Divides in Infectious Disease",
Borer ET, Antonovics J, Kinkel LL, Hudson PJ, Daszak P, Ferrari MJ, Garrett KA, Parrish CR, Read AF, Rizzo DM. Department of Ecology, Evolution, and Behavior, University of Minnesota, St. Paul, MN, 55108, USA, Ecohealth. 2011 Nov 16. [Epub ahead of print]

purpose	original [98 words]
	Pathogens traverse disciplinary and taxonomic boundaries, yet infectious disease research occurs in many separate disciplines including plant pathology, veterinary and human medicine, and ecological and evolutionary sciences.
	These disciplines have different traditions, goals, and terminology, creating gaps in communication.
	Bridging these disciplinary and taxonomic gaps promises novel insights and important synergistic advances in control of infectious disease.
	An approach integrated across the plant-animal divide would advance our understanding of disease by quantifying critical processes including transmission, community interactions, pathogen evolution, and complexity at multiple spatial and temporal scales.
	These advances require more substantial investment in basic disease research.

Here is, for each abstract my Purpose version, followed by my Polished version. Study them and compare them to your versions. Did I lose the story?

Purpose

purpose	original [151 words]
context	Several data mining algorithms use iterative optimization methods for learning predictive models.
problem	It is not easy to determine upfront which optimization method will perform best or converge fast for such tasks.
what we did	In this paper, we analyze Meta Algorithms (MAs) which work by adaptively combining iterates from a pool of base optimization algorithms.
result1	We show that the performance of MAs are competitive with the best convex combination of the iterates from the base algorithms for online as well as batch convex optimization problems.
evidence	We illustrate the effectiveness of MAs on the problem of portfolio selection in the stock market and use several existing ideas for portfolio selection as base algorithms.
result2	Using daily S&P500 data for the past 21 years and a benchmark NYSE dataset, we show that MAs outperform existing portfolio selection algorithms with provable guarantees by several orders of magnitude, and match the performance of the best heuristics in the pool.

Polish

purpose	revised [94 words]
context problem 22	To learn predictive models, data mining algorithms using iterative optimization methods cannot determine beforehand which method will perform best or converge fast.
what we did 15	We analyze Meta Algorithms (MAs) which adaptively combine iterates from a base optimization algorithms pool.
result1 19	MAs, for online and batch convex optimization problems, are competitive with the best convex combination of base algorithms iterates.
result2/evidence 38	Using known portfolio selection ideas as base algorithms and two decades of daily S&P500 and a benchmark NYSE dataset, MAs select stock market portfolios effectively, strongly outperforming existing algorithms and matching the performance of the best pool heuristics.

Purpose

purpose	original [98 words]
context	Pathogens traverse disciplinary and taxonomic boundaries, yet infectious disease research occurs in many separate disciplines including plant pathology, veterinary and human medicine, and ecological and evolutionary sciences.
problem	These disciplines have different traditions, goals, and terminology, creating gaps in communication.
potential impact	Bridging these disciplinary and taxonomic gaps promises novel insights and important synergistic advances in control of infectious disease.
detail	An approach integrated across the plant-animal divide would advance our understanding of disease by quantifying critical processes including transmission, community interactions, pathogen evolution, and complexity at multiple spatial and temporal scales.
conclusion	These advances require more substantial investment in basic disease research.

Polish

purpose	first revision [60 words]
context	Pathogens traverse disciplinary and taxonomic boundaries, but infectious disease research does not.
problem	Differing traditions, goals, and terminology create communication gaps.
potential impact	Bridging gaps promises novel insights and synergistic advances in controlling infectious disease.
detail	Integrating across the plant-animal divide would quantify, at multiple spatial and temporal scales, critical processes including transmission, community interactions, and pathogen evolution.
conclusion	However, bridging gaps requires substantial investment.

Which of the 50 word versions do you prefer?

purpose	second revision [50 words]
context	Pathogens traverse disciplines and taxonomies, but infectious disease research does not.
problem	Differing traditions and terminology create communication gaps.
potential impact	Bridging gaps promises insights and synergistic advances.
detail	Integrating across the plant-animal divide would quantify, at multiple levels, critical processes including transmission, community interactions, and pathogen evolution.
xonclusion	However, bridging gaps requires substantial investment.

purpose	alternate second revision [50 words]
context/problem	While pathogens traverse disciplines and taxonomies, isolated infectious disease research creates communication gaps
detail	Traditions and terminology differ in plant pathology, veterinary and human medicine, and ecological and evolutionary sciences.
potential impac	Bridging gaps would quantify, at multiple levels, critical transmission, community interactions, and pathogen evolution processes.
conclusion	However, bridging gaps requires substantial investment.

5. Abstract Storylines.

Below are abstract storylines created by me or graduate students (mostly in Computer Science and Environmental Engineering) attending workshops. Each storyline is a left hand column of a table rotated up). Each is a list of semicolon delimited sentence purposes.

Work: As you create storylines, keep them together so you can study them. Aim to acquire a feel for what makes a good abstract.

context; the need / competition; approach; method, results

context; background, purpose / reason; approach, findings

Crux; Supplied; Supplied

context; prior art; contribution; method name; approach; case study; challenge; contribution; results

Context; BUT; Approach; Method; Relevance of our work; What is Special; Results and Impact

Approach; Purpose; Approach difficulties; Proof that approach is viable; Contribution to field; Novelty Justification; Secondary outcome of note

motivation behind recsys*; recsys successful approaches; assumption of successful approaches; weakness of assumption (opportunity); weakness of assumption; weakness of assumption; weakness of assumption; problem & approach; method; experimental results; real-life application results

context; approach; method; outcome; evidence; results; ways to use

Context; Opportunity; Approach; Why use; Method; Kind of evidence

context; problem; purpose; terms; definition; definition; scope; caveat

Contribution; Problem/ Previous work; Purpose/Why; Method; Results (?)

Context; Overview; Motivation/method partly; Contribution; Method/ what was done; Results

Context; Approach; Results; Method; Opportunity; Method; Results; Kind of Evidence; Further Applications

Context; BUT; Approach; Method; Relevance of our work; What is Special; Results and Impact

Context; Opportunity; Approach; Method; Crux; Why use; Kind of evidence & results

Background; Existing protocols; Problems; Concern #1; Concern #2; Reason; Approach; Strength; Theoretical result; Simulation result

Introduce underlying technologies; What underlying technologies offer; broad approach; specifics to approach; What will be discussed; introduce findings; summary of specific findings; applications of findings

Context; Opportunity; Motivation; Crux; Approach; Method; Result

Problem ; Motivation ; Work overview ; Employed techniques ; Elaborate on work ; Results overview ; Datasets ; Results ; Conclusion

Overview ; Problem ; Motivation ; Contribution ; Contribution ; Contribution ; Evaluation metrics ; Results

Context; Opportunity; Approach; Method; Crux; Results; impact

Context; Opportunity; Approach; Results; More results; Efficiency

background; Additional information; approach; Method and result; Further experiment result; Further result in detail; Analyze

context; but; but; opportunity; crux; method; evidence/impact

problem statement; running example; application domain; computational challenges; related work; approach; novelty; validation; better

research question / sampling framework; variables tested; result detail; result detail; immediate conclusion; larger impact; ; ; ; ; ;

Approach; Results 1; Results 2 ; Results 3; Results 4/ Implication; Potential Impact; ; ; ; ; ;

what we did; approach; approach; approach; why important; results; results; results; impact; ; ;

Context; Results; Results; Results; Results; Results; Results; Broad Results; Broad Results; Pur-pose; ;

context ; problem; opportunity; method; method; /results; results; results; impact; ; ;

Opportunity; Research Focus; Method; Analysis; Analysis; Hypo-thesis; Hypo-thesis; Results; Conclusion; Anal-ysis; Take home message; Take home message

Context ; Opportunity ; Opportunity ; Purpose of Article; Approach ; Challenges ; ; ; ; ; ;

context; variables; methods; Key result; Results; Results; Conclusion; impact; impact; impor-tance; contra-diction;

Background; Approach; Method; Results of research; Impact of results; Verdict; ; ; ; ; ;

problem; approach; result; bigger result; conclusion; impact; ; ; ; ; ;

Purpose; Methods; Methods; Methods; Results; Results; Results; Results; ; ; ;

Context; Context; Method; Method; Results; Results; Results; Results; Results; ; ;