Getting started

with PMark

Paul Anderson

<dcspaul@ed.ac.uk>

Software & Documentation

http://homepages.inf.ed.ac.uk/dcspaul/pmark

PMark

- Uncompromisingly criteria-based mark calculator !
- Not the tool to use if you want to "add up marks", "accumulate credit", ...
- Research: how far can we get with just criteria-based rules ?
  (this has lots of benefits)

This session

- I am really interested in discussing the applications & pedagogy
- But this is a hands-on session about the practical usage of the tool
- I'm particularly interested in seeing what people find difficult/awkward
- Please tell me when something doesn't make sense!

The interface

- PMark has a web interface, but the flexibility that the tool offers comes from the ability to express the requirements in a text "language"
- So we will mostly be talking about the PMark "language"

Resources

PMark software

- Is freely available and open source
- Command line version in use for several years, but requires Unix or Mac
- Web server version is new - interface less well tested

PMark service

- A version of the server for University staff running in Informatics
- "Best effort" support! Talk to me if you want to use it for "production"

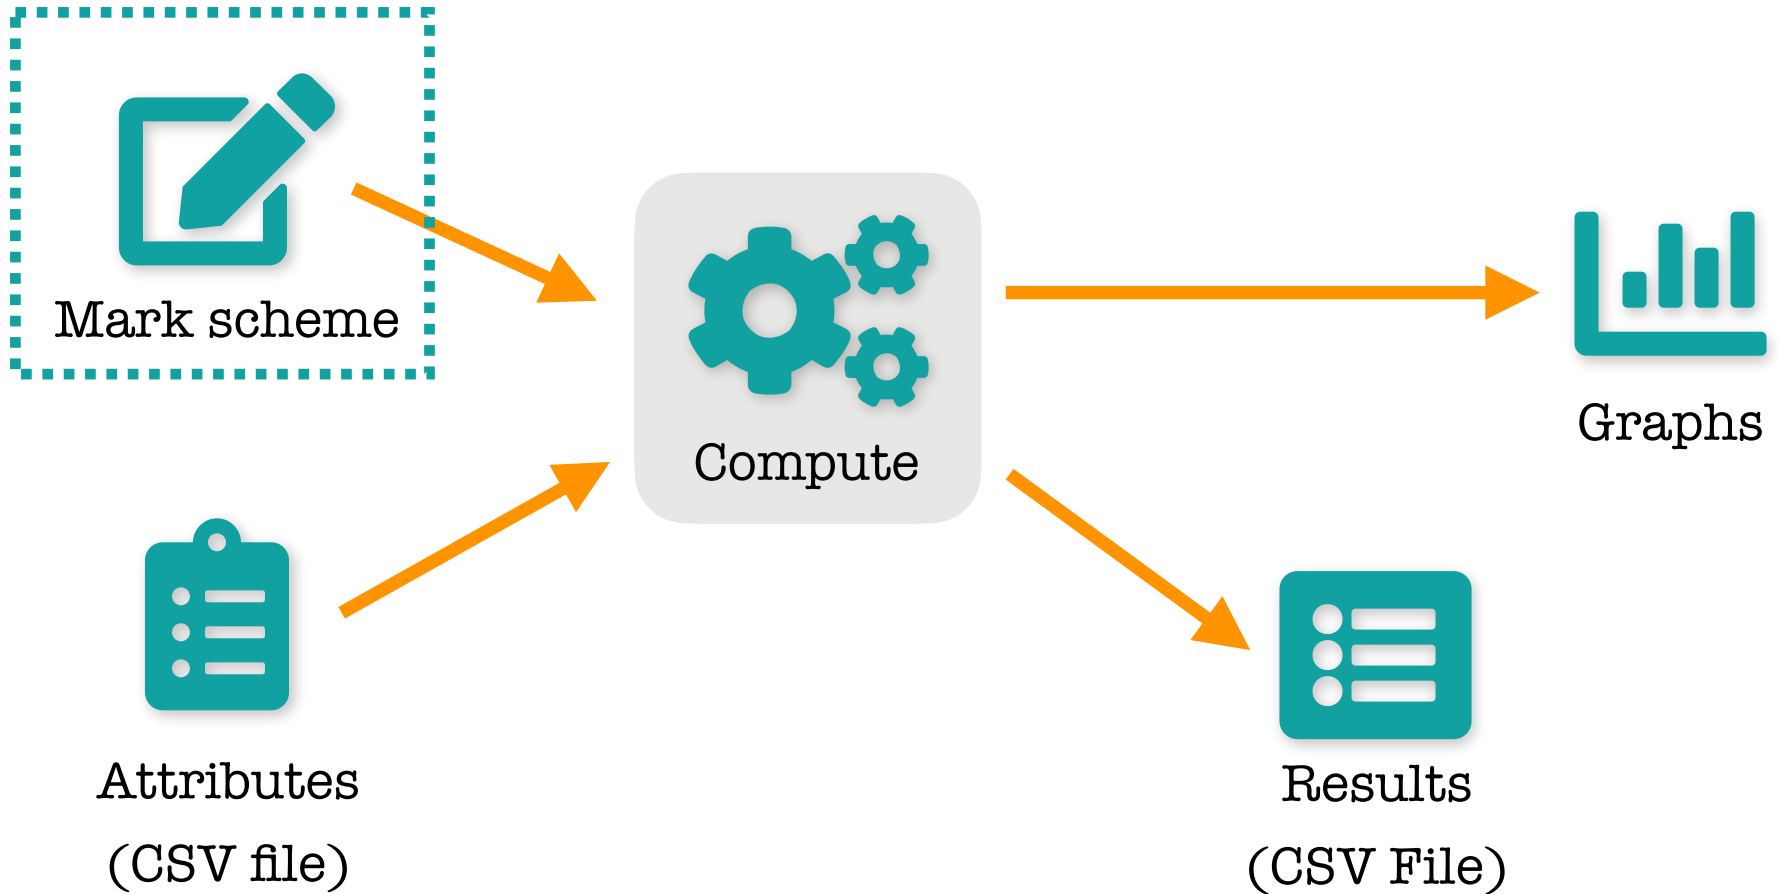https://ease.sweb.inf.ed.ac.uk/dcspaul/pmark/master/pmark.cgi/server

Documentation

- Introduction document, reference document
- Beginnings of something discussing the practical applications & pedagogy
- Talk videos / slides & basic demonstration videos

http://homepages.inf.ed.ac.uk/dcspaul/pmark

PMark Overview

Mark scheme

Compute

Graphs

Attributes

(CSV file)

Results

(CSV File)

id	content	spelling
Alexander	good	good
David	good	bad
Lucy	bad	good
Silas	bad	bad

id	final
Alexander	pass
David	fail
Lucy	pass
Silas	fail

A PMark data file

The file is a CSV file
- created with Excel/Notepad
- or with PMark forms

You must have a (lower case) "id" column
with (unique) student "identifiers"
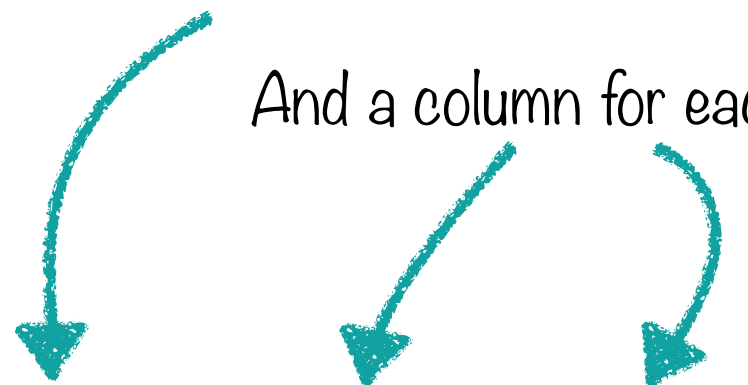
The "identifiers" must be short names
- with only letters, numbers or hyphens
- starting with a letter
- the case is significant

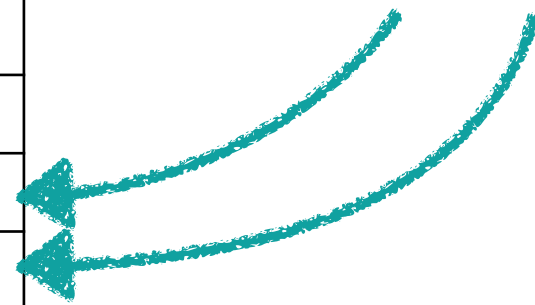And a column for each attribute (the names are "identifiers")

The values can be "identifiers"
or (whole) numbers

id	content	spelling
Alexander	good	good
David	good	bad
Lucy	bad	good
Silas	bad	bad

A PMark mark scheme

The mark scheme is a plain text file
- created in the web interface
- or with an editor - eg. Notepad
- it usually has a ".pmark" file type

It has a number of sections, each headed by the section type in (square) brackets ...

This lists the different types of attribute that we might have: eg. pass/fail, likert, percentage, ...

This lists the criteria (attributes) and their corresponding types

This lists the results and their corresponding types

This gives the rules relating the results to the attributes

[types]

...

[attributes]

...

[results]

...

[rules]

...

PMark types

We only have one type here
It has two possible values:
"bad" and "good"

```
[types]
mark: [ bad, good ]
```

You can call the type whatever you
like (as long as it is an "identifier")
I have called it "mark"
You could call it 🧁

The punctuation is important!
The colon and the brackets

The order is important!
"good" is "better" than "bad"

id	content	spelling
Alexander	good	good
David	good	bad
Lucy	bad	good
Silas	bad	bad

We have two attributes and they both have the same type

[types]

mark: [bad, good]

[attributes]

content: mark

spelling: mark

id	content	spelling
Alexander	good	good
David	good	bad
Lucy	bad	good
Silas	bad	bad

A real mark scheme may have dozens of attributes with several different types

Creating some data

We now have enough of a mark scheme to describe the data
PMark can automatically generate some matching (random) data ..

Enter the mark scheme

- Click on the Scheme icon
- Click the New button
- Give the scheme a (short) name
- Type in the scheme
- Click Save
- If there are any typos,
  correct them and click Save again

Generate some data

- Click Random
- Give the data a (short) name

```
[types]
mark: [ bad, good ]

[attributes]
content: mark
spelling: mark
```

When you have"real" marking data, you
can upload it from your CSV file
or enter it using PMark forms

We only have one result
which we have called "final"
It it specified in the same
way as the attributes ...

It has a "type" which we
have called "grade"
You could call this anything
you like

The grade has two values
"fail" or "pass"
(be careful of the order!)

[types]
grade: [fail, pass]

[results]
final: grade

id	final
Alexander	pass
David	fail
Lucy	pass
Silas	fail

The rules tell us how to compute the result, from the values of the attributes

We have one very simple rule: the mark is going to be a pass (only) if the content is "good"

If none of the rules are satisfied, then the result grade will be the lowest possible ie. a "fail"

The spelling will be ignored for now

```
[types]
mark: [ bad, good ]
grade: [ fail, pass ]


[attributes]
content: mark
spelling: mark


[results]
final: grade


[rules]
pass: content=good
```
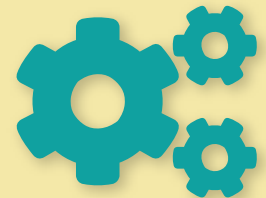
This scheme is very simple! But it is enough for PMark to compute a result

Update the mark scheme

- Edit the mark scheme to look like the version on the right
- Click Save
- If there are any typos, correct them and click Save again

Compute the result

- Click Compute

PMark will compute the result and you should be able to see that (only) those students with good content have passed

```
[types]
mark: [ bad, good ]
grade: [ fail, pass ]


[attributes]
content: mark
spelling: mark


[results]
final: grade


[rules]
pass: content=good
```

A basic rules says that an attribute must have <u>at least</u> a certain value

So this rule doesn't say quite what it first looks like!

[types]
mark: [terrible, bad, good]

[rules]
pass: content = bad

[types]
grade: [fail, pass, distinction]

[rules]
distinction: content = good
    and spelling = good

A rule can require several criteria (connect them with "and")

You can also require just one criteria or the other (connect them with "or")

One rule can depend on another rule

We almost almost always want the higher grades to depend on the lower ones in this way ...

[rules]

distinction: pass

    and spelling = good

[rules]

distinction: most of {

    spelling = good

    content = good

    ....

}

We can also ask for ...
one of { ... } , some of { ... } , most of { ... }
all of { ... } , 7 of { ... } , all but one of { ... }
all but 3 of { ... }

Note the punctuation (curly brackets)

Finally: we can combine any of these to make the criteria as complex as we need!

An exercise

Try extending the mark scheme

- Add a "distinction" to the grades
- Add a new rule to award a distinction if the spelling is good as well as the content
- Run this on your random data and check a few results

Some other things to try ...

- Add an "exceptional" to the grades
- Add an extra attribute for originality with a value of 0-5
- Invent a new rule for the exceptional grade

You will need to add a new type for the numeric attribute

You will need to generate a new data file when you have added a new attribute

```
[types]
mark: [ bad, good ]
grade: [ fail, pass ]

[attributes]
content: mark
spelling: mark

[results]
final: grade

[rules]
pass: content=good
```

Numeric results

It is easy to create a numeric type

And there is a shorthand (two dots) if the type has a big range

There is a limit (about 150) on the range

```
[types]
likert: [ 0,1,2,3,4,5 ]
percentage: [ 0..100 ]
```

```
[types]
cms: [ 0..100
  H=0, G=10, F=20, E=30
  D=40, C=50, B=60
  A3=70, A2=80, A1=90
]
```

But we can't use numbers for the names of the rules, so ...
If we want a numeric type for the result, we need to attach names to the values corresponding to the rules

```
[rules]
B: content = good
```

```
[types]
likert: [ 0..5 ]
mark: [ bad, good ]
cms: [ 0..100
  pass=50
  distinction=70
  exceptional=80
]

[attributes]
content: likert
spelling: mark
originality: likert
```

```
[results]
final: cms

[rules]
pass: content=2
distinction: pass and spelling=good
exceptional: distinction and originality=3
```

Notice that there are 101 possible values for the result, but only 3 rules ..

If the result type has values with no corresponding rules, PMark will "interpolate" the values inbetween depending on how well the rules have been met
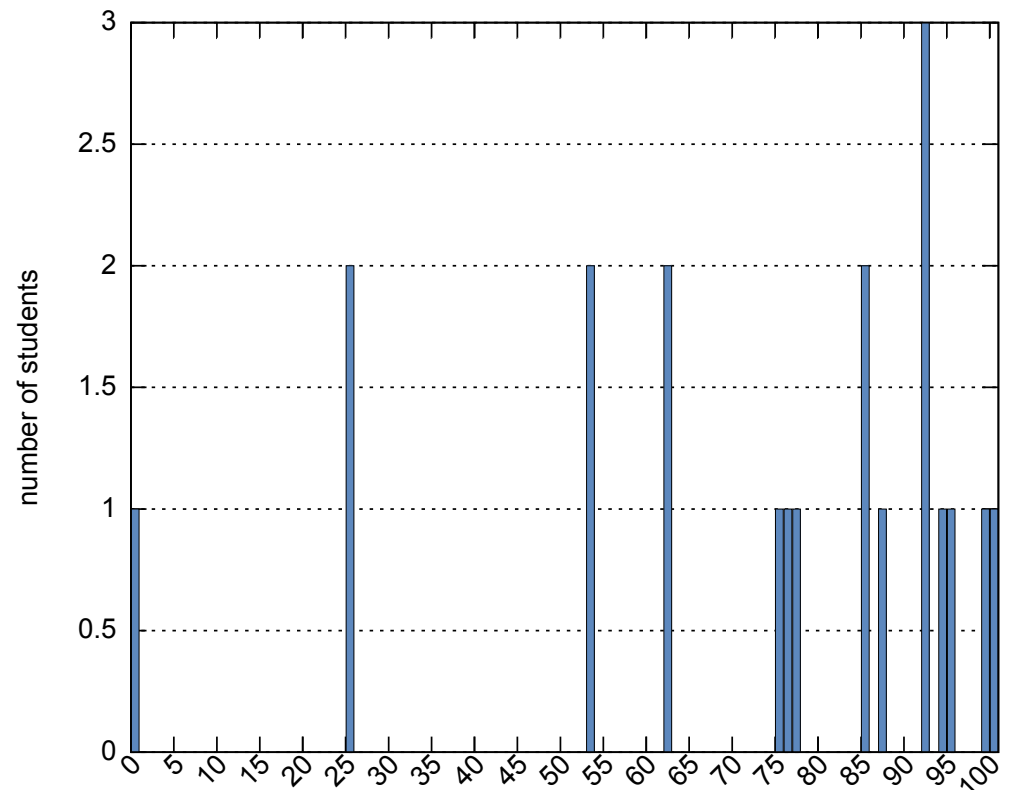
id	content	originality	spelling	final
Harper	3	4	good	92
Heather	2	3	good	85

Heather and Harper are both "excellent", but Harper has done slightly better …

The interpolation is not perfect, but it does generate plausible values between grade points

But: The interpolation never violates the rules!

Hashtags

In practice, typical schemes have lots of attributes

To make it easy to refer to groups of them, we can attach (one or more) "hashtags" to an attribute

[rules]

pass: most #abstract = 3

[attributes]
wordcount: likert #abstract
findings: likert #abstract
conclusions: likert #abstract
background: likert #abstract
methodology: likert #abstract
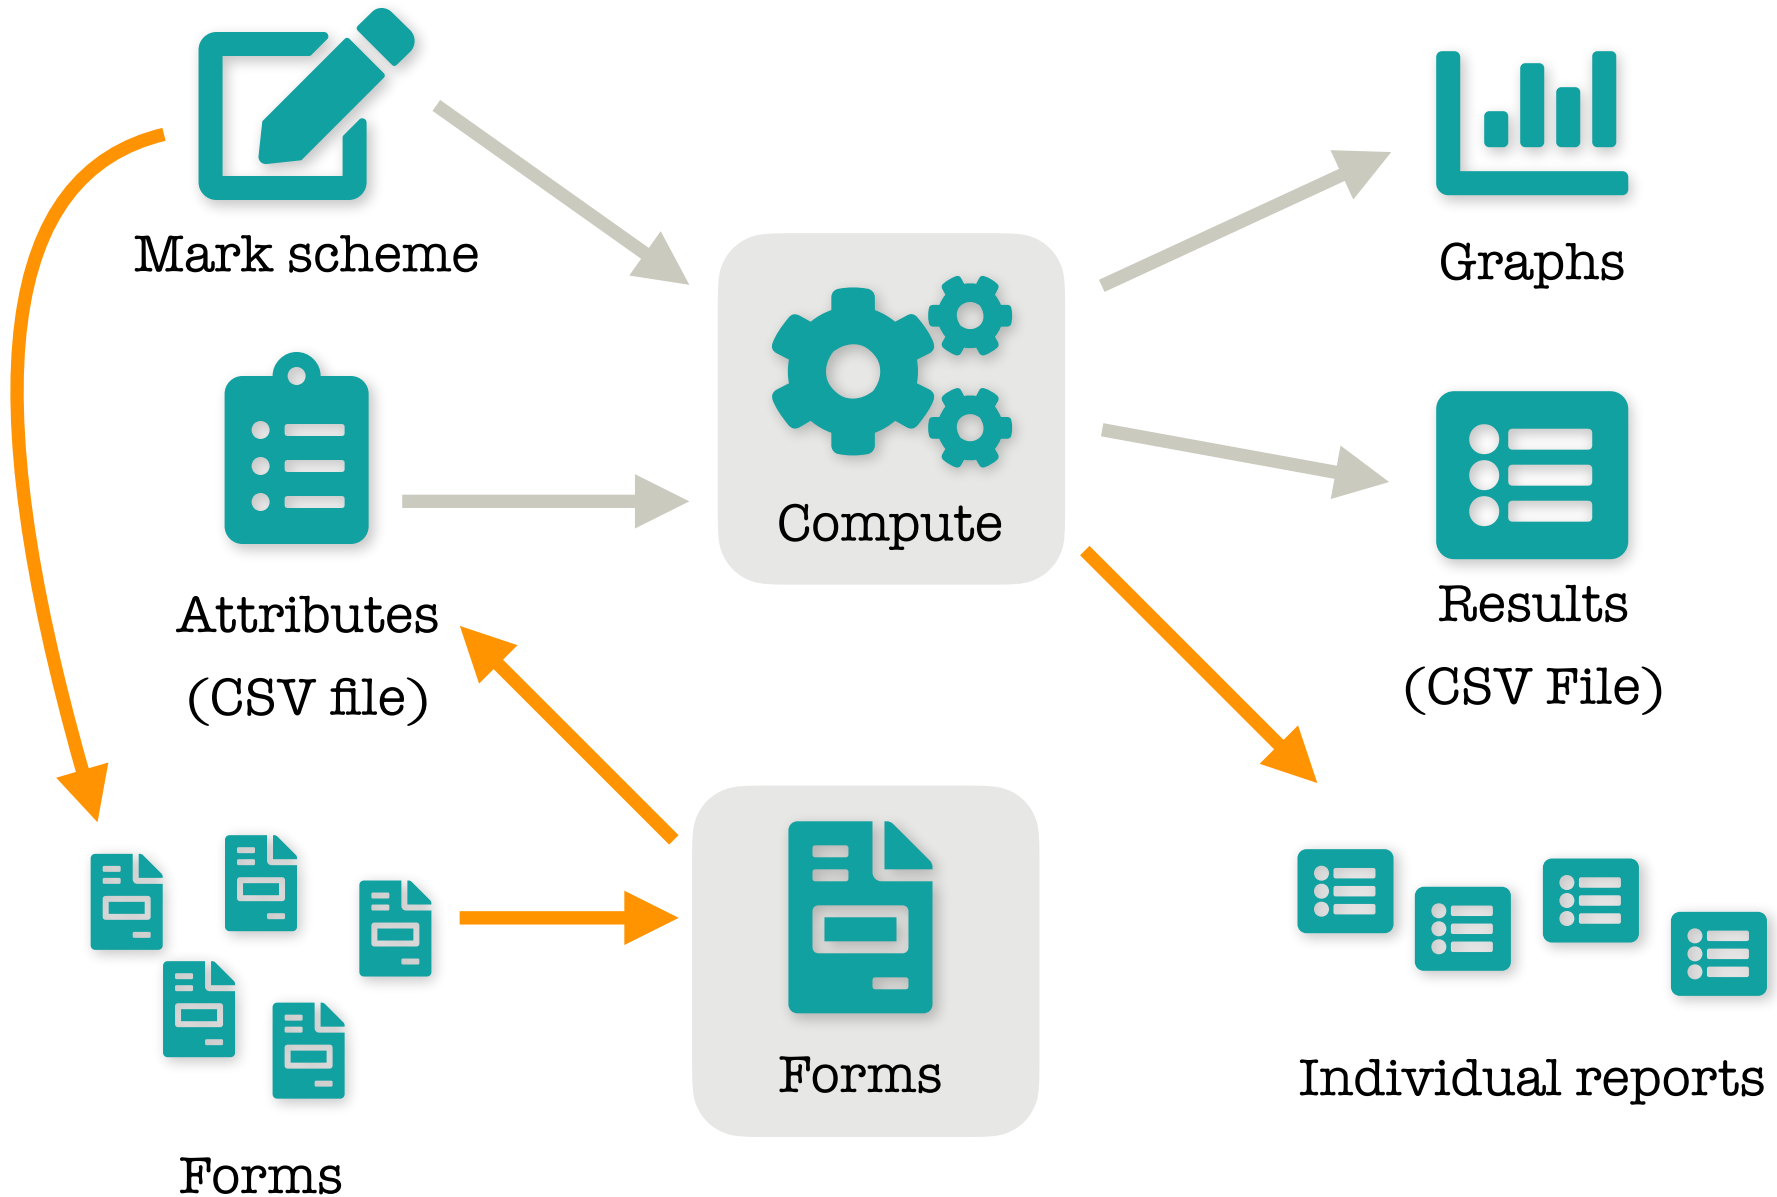question: likert #abstract
...

This is common way of expressing a degree of leniency ...

[rules]

pass: most #abstract = 4 and all #abstract = 3

What else can we do ?

Mark scheme

Attributes
(CSV file)

Forms

Compute

Forms

Graphs

Results
(CSV File)

Individual reports

group10a
group10b
group1a •
group1b •
group2a
group2b
group3a
group3b
group4a
group4b •
group5a
group5b
group6a •
group6b
group7a
group7b
group8a
group8b
group9a
group9b

Show All   Hide All

Cancel   Save

Report

Title & affiliations

Does the title describe the report in a clear & concise way ?	not-really
Are the authors' surnames listed in alphabetical order ?	yes

Abstract

Is the word count between 100-200 words ?	no
Are the main findings clearly summarized ?	yes
Are the main conclusions clearly summarized ?	no
Is the background to the research presented ?	yes-but
Is the methodology used to answer the research question clearly described ?	?
Is the research question clearly stated ?	?

Body

Is the introduction different from the abstract ?	?
Is the background to the research presented in sufficient depth to	?

OK   new form for 'group2a'

Paul Anderson <dcspaul>   PMark 4.4

🎗 Marks for Delaney report: 64%

Delaney achieved a pass (40) for the report-mark.

Suggestions for a higher grade ...

For a distinction (70) ...

1. The background to the research must be presented clearly and completely.
2. The sources should be clearly credited within the body of the work.
3. The main conclusions must be completely & clearly summarized.
4. The methodology to the research must be described clearly and completely.
5. The research question must be stated clearly and completely.
6. There must be a clear demonstration of critical analysis of sources.
7. The research question must be presented in depth.
8. The conclusions should be clearly based on the findings.
9. The limitations of the findings must be clearly described.
10. The methodology question must be presented in depth.
11. The sources must be well integrated into the report.
12. Full citations should be given at the end in a reference list.
13. All sources should be pre-dominantly peer-reviewed (or an indication provided as to why not).
14. The title must describe the report clearly and completely.

Where next ?

More examples

- Copies of the examples from the introduction document are available from the Scheme menu (you can generate test data for these and then run/copy/edit them)
- Trying to implement some assessment ideas of your own is probably the best way to understand PMark
- Developing the rules usually forces people to think more deeply about their assessments and this takes time "up front"

Support

- Please talk to me if you are thinking of using this on a "real" assessment
- I will be an honorary fellow from September & I am interested in continuing this work, so I would be happy to talk to anyone who is interested in following up
- I am really interested in your (honest!) feedback