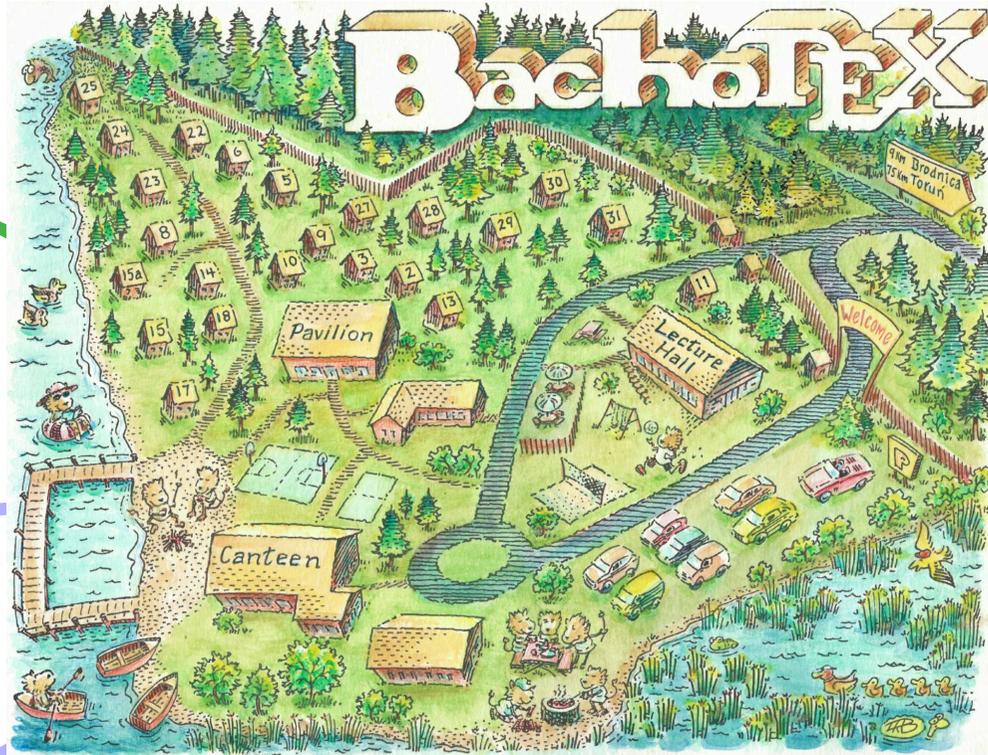


## Alice in Bachotek

Frank Mittelbach  
May 2016



## Alice in Bachotek

Frank Mittelbach  
May 2016



Conclusions



Roadwork



The tale of the  
tail of bugs



Introduction



The pagination  
problem



The pagination  
framework



First results



What's this all about?

Example: Typesetting Alice



So why do we do all this?

Excurs: Typesetting TLC2

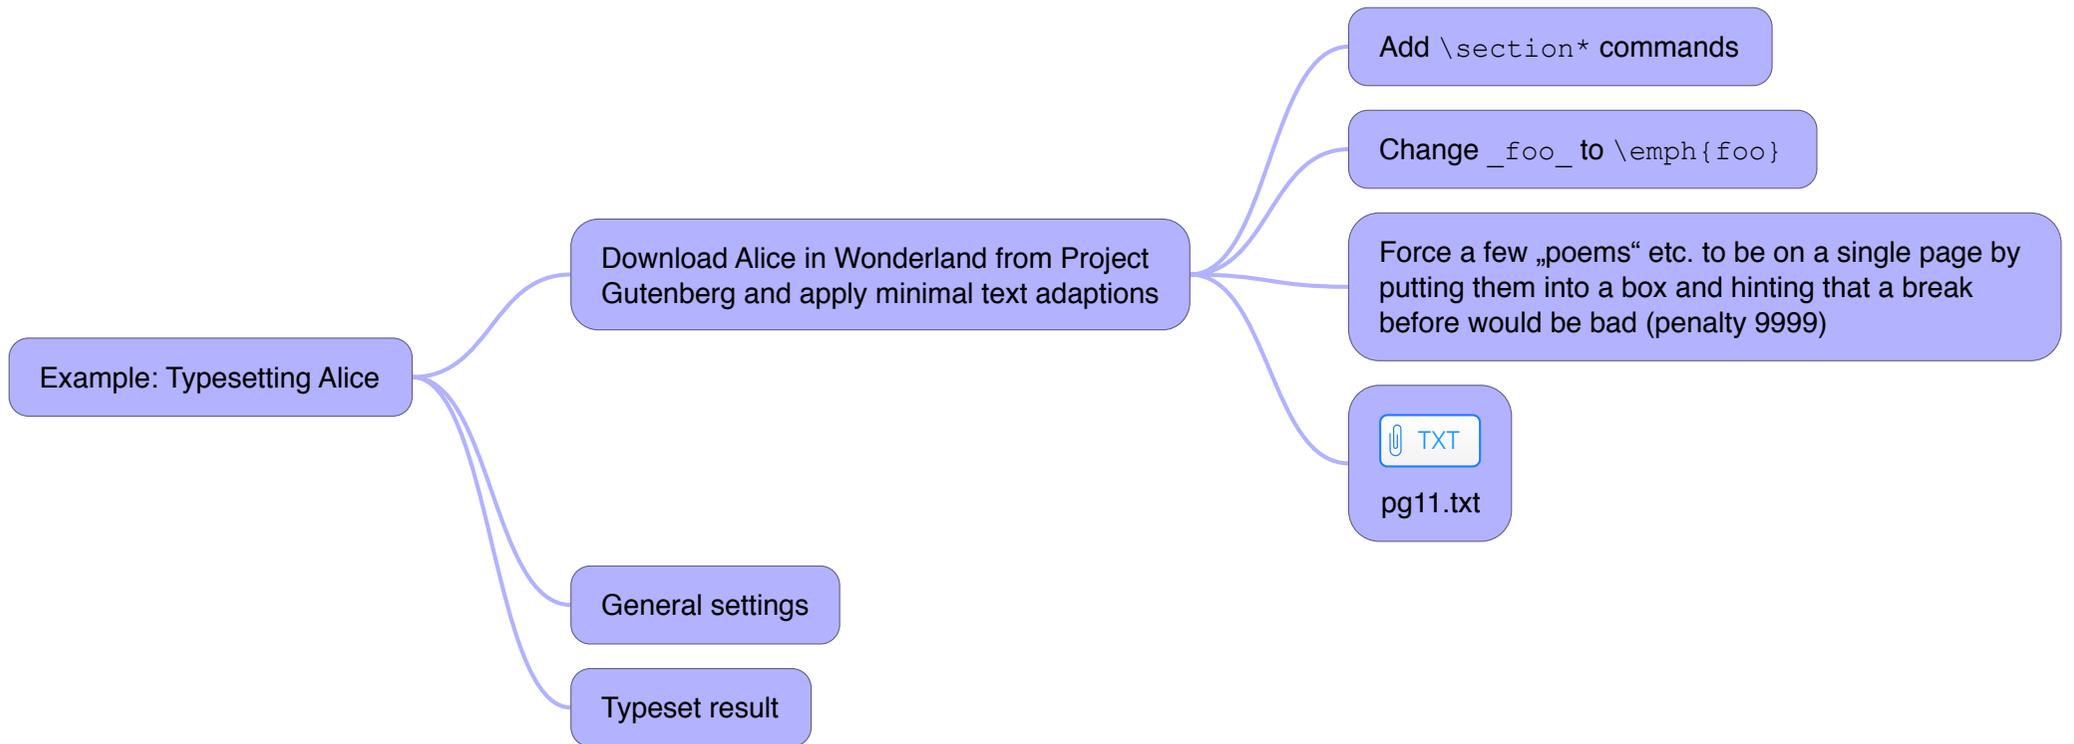
Example: Typesetting Alice

```
graph LR; A[Example: Typesetting Alice] --- B[Download Alice in Wonderland from Project Gutenberg and apply minimal text adaptations]; A --- C[General settings]; A --- D[Typeset result];
```

Download Alice in Wonderland from Project Gutenberg and apply minimal text adaptations

General settings

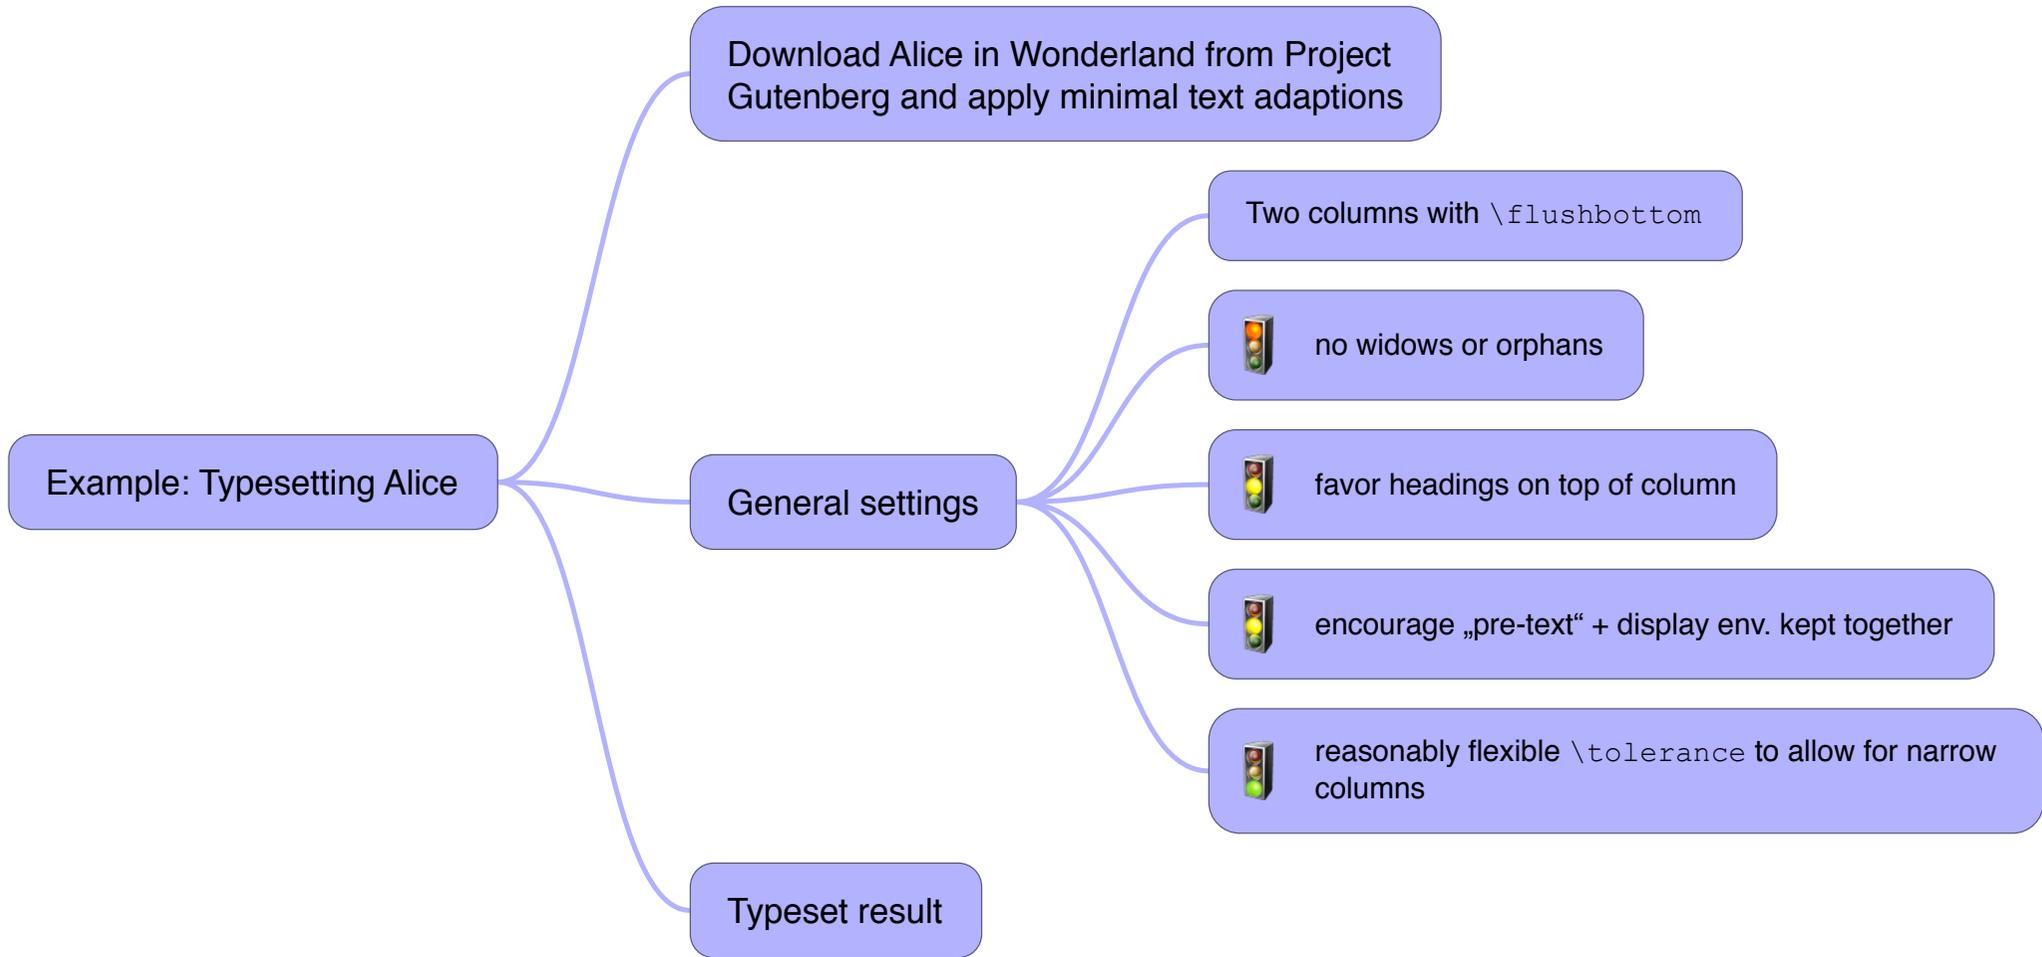
Typeset result





TXT

**pg11.txt**



Example: Typesetting Alice

Download Alice in Wonderland from Project Gutenberg and apply minimal text adaptations

General settings

Two columns with `\flushbottom`

 no widows or orphans

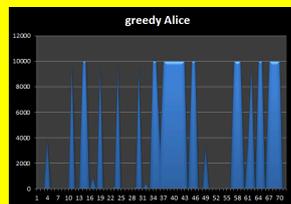
 favor headings on top of column

 encourage „pre-text“ + display env. kept together

 reasonably flexible `\tolerance` to allow for narrow columns

Typeset result

Typeset result



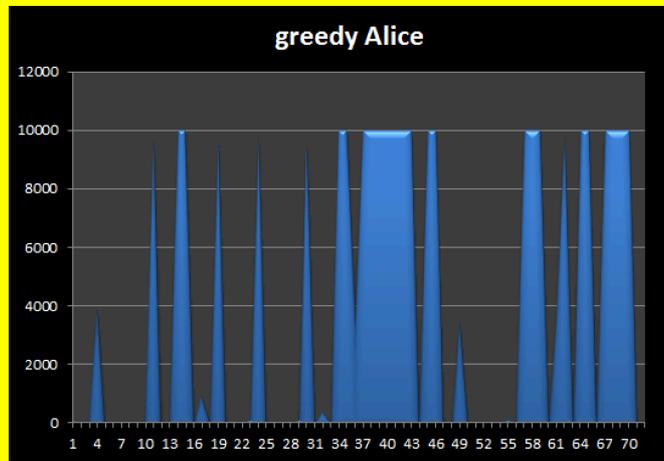
Medium disaster



Typeset result by standard LaTeX



Consequences



Medium desaster



Typeset result by standard LaTeX

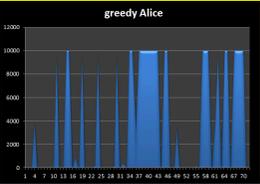
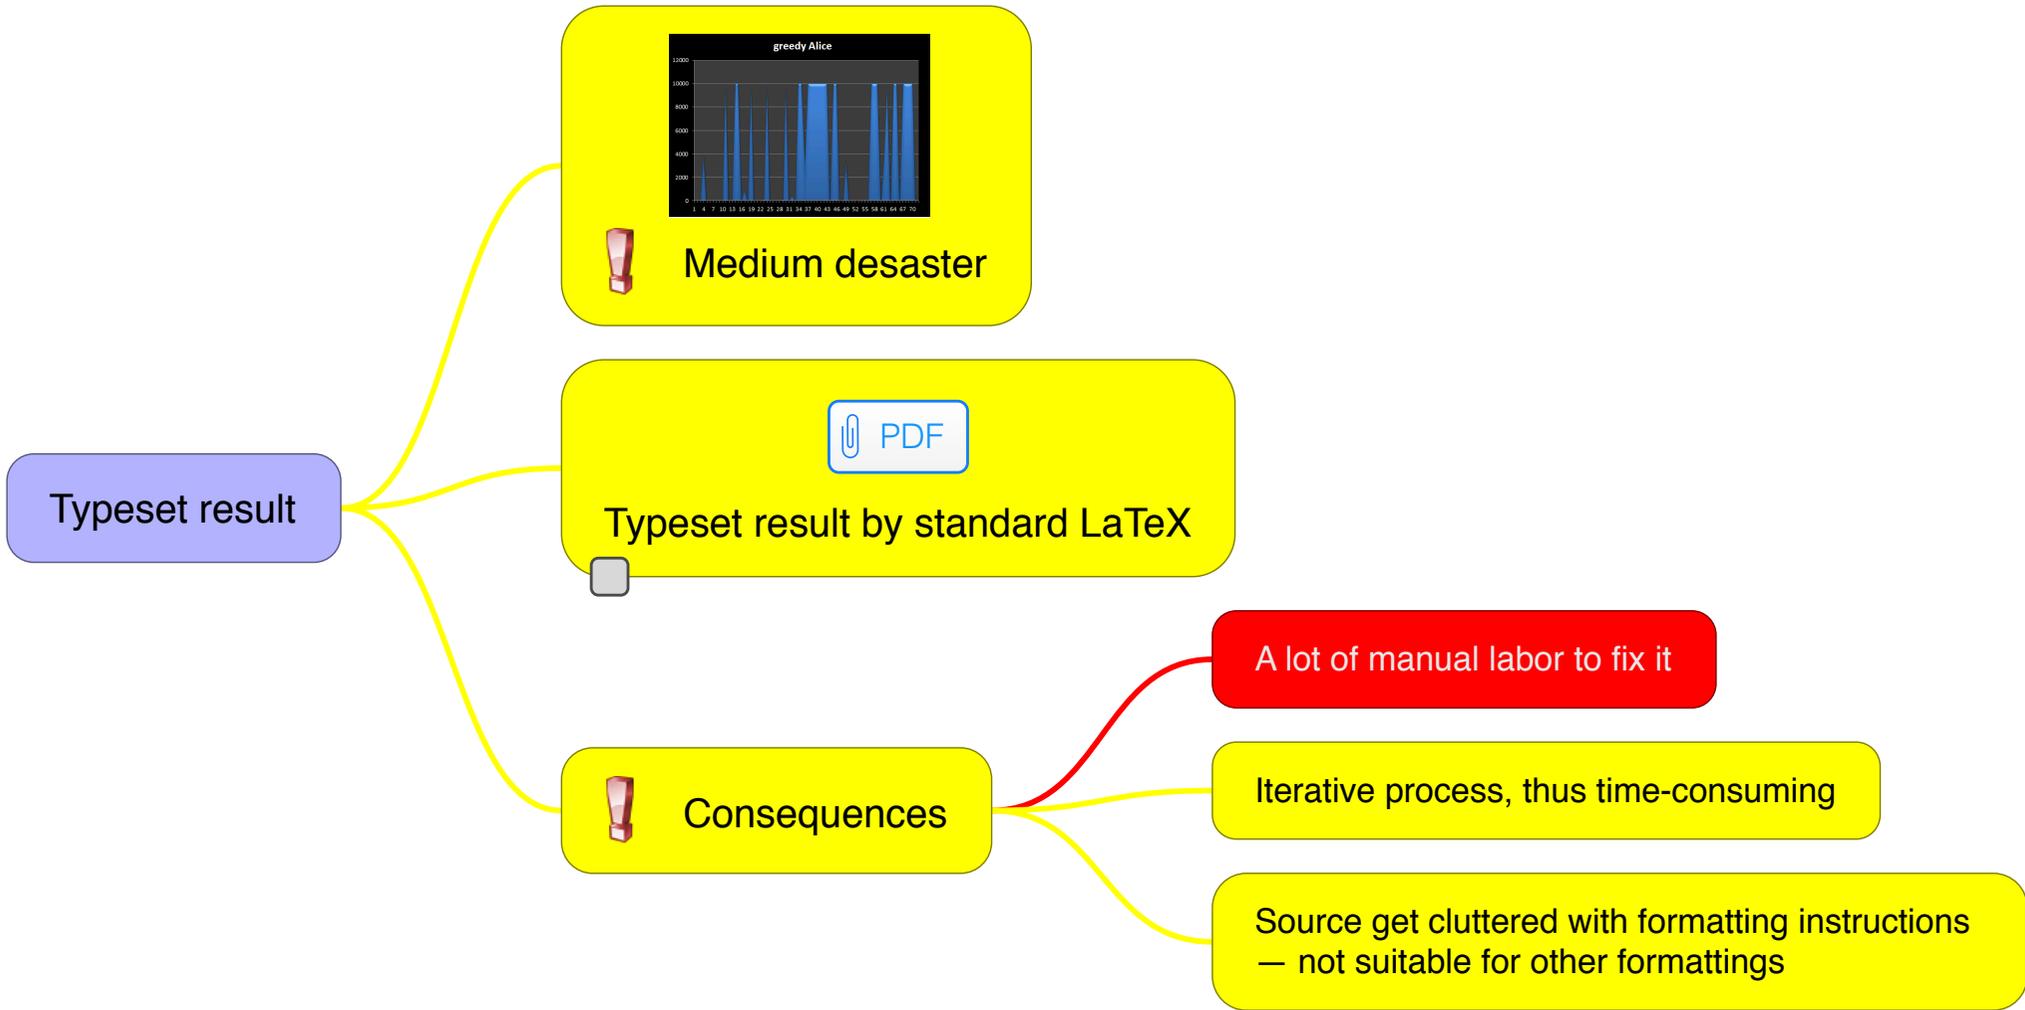
Page 6 — large space in front of heading

Page 7 — second column 2/3 empty  
(because of mouse tail following)

Page 10 — 2 lines missing at the bottom  
because of 3-line paragraph following

Page 12 — 1 line missing at the bottom  
because of poem following + **bad break!**

On some of the later pages LaTeX ran both  
columns short so visually the results are fine



Medium disaster



Typeset result by standard LaTeX



Consequences

A lot of manual labor to fix it

Iterative process, thus time-consuming

Source get cluttered with formatting instructions – not suitable for other formattings



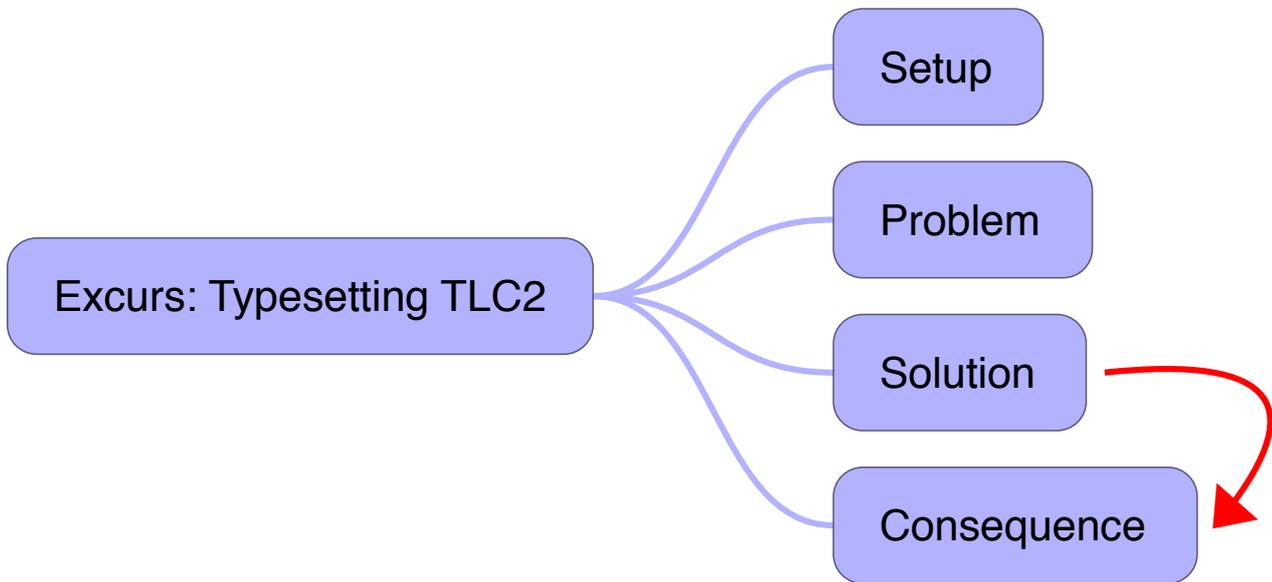
What's this all about?

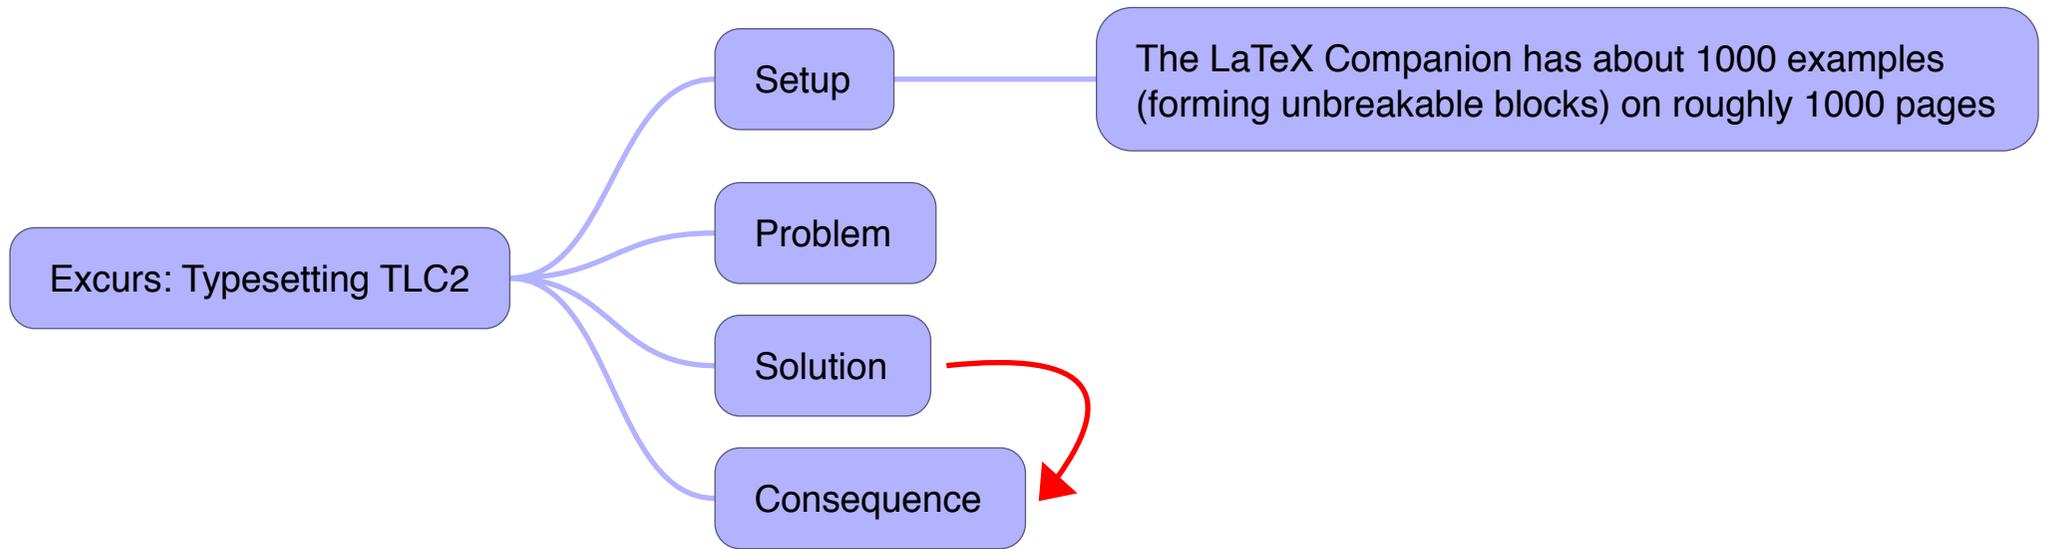
Example: Typesetting Alice

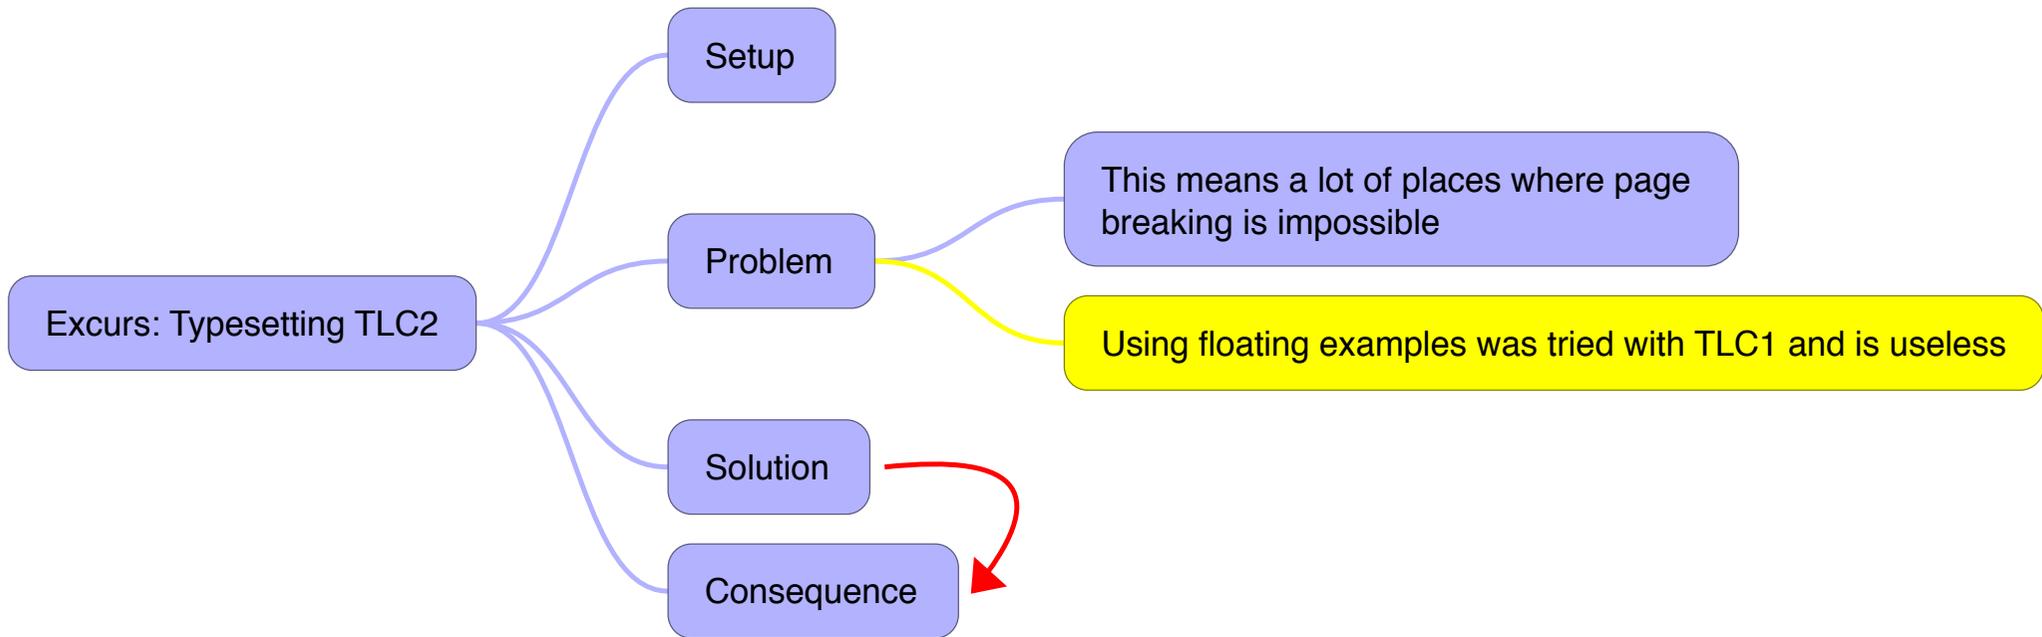


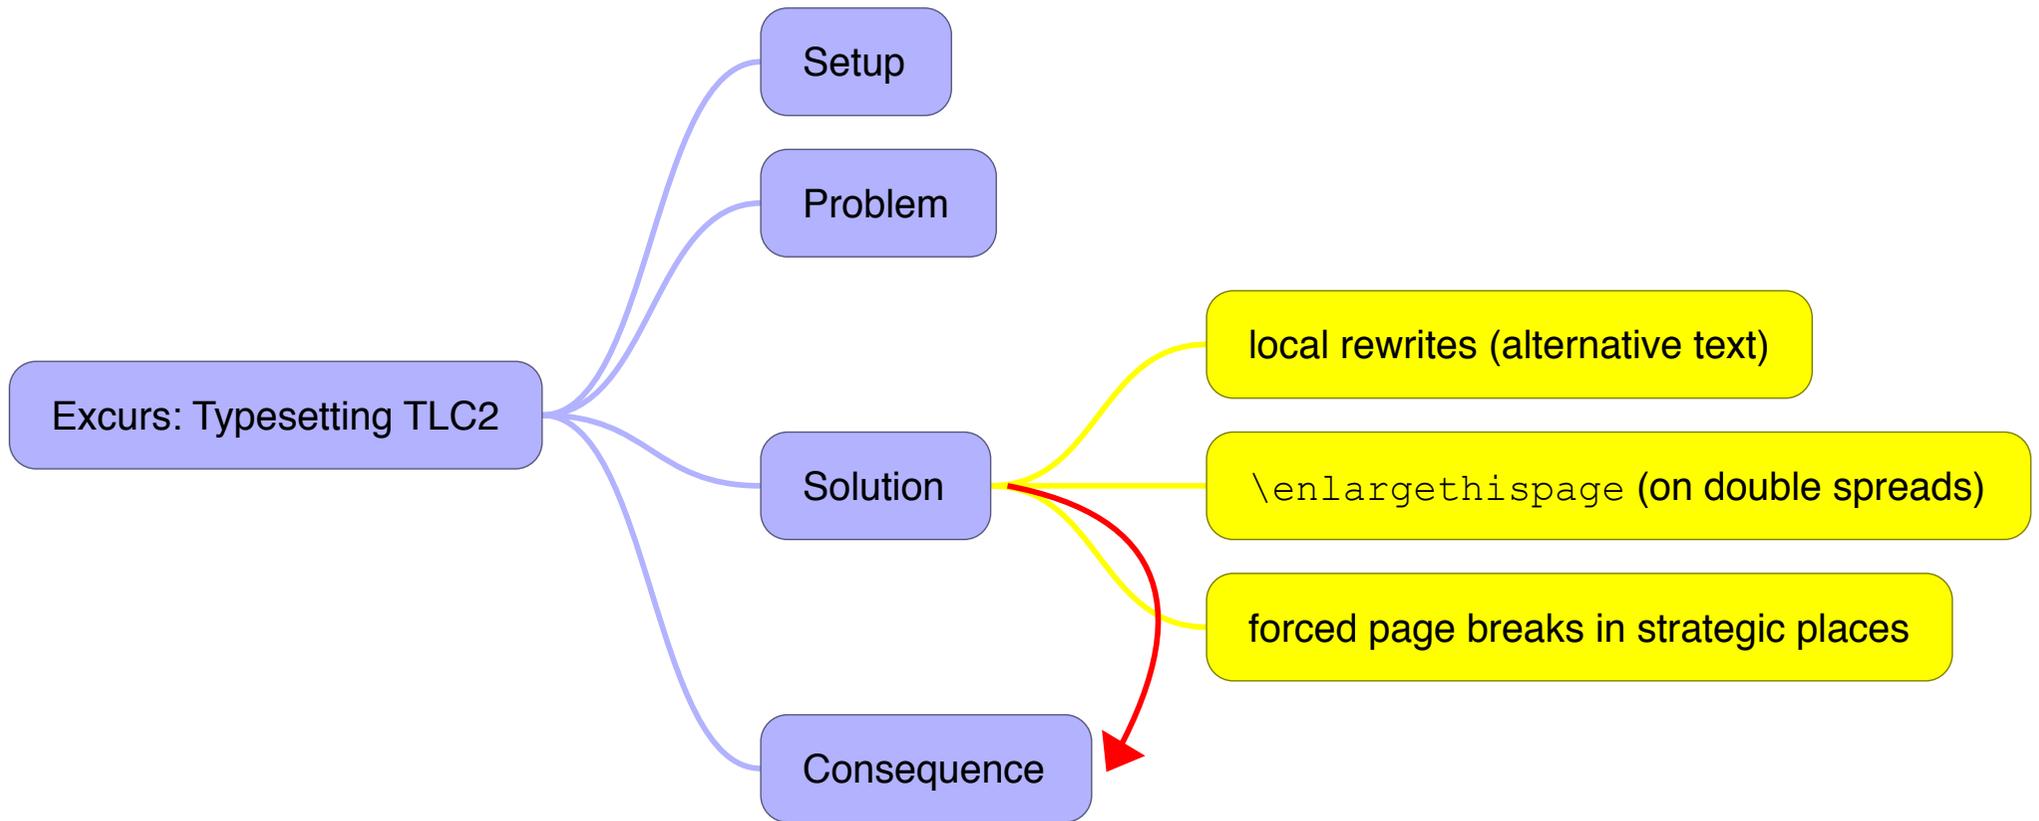
So why do we do all this?

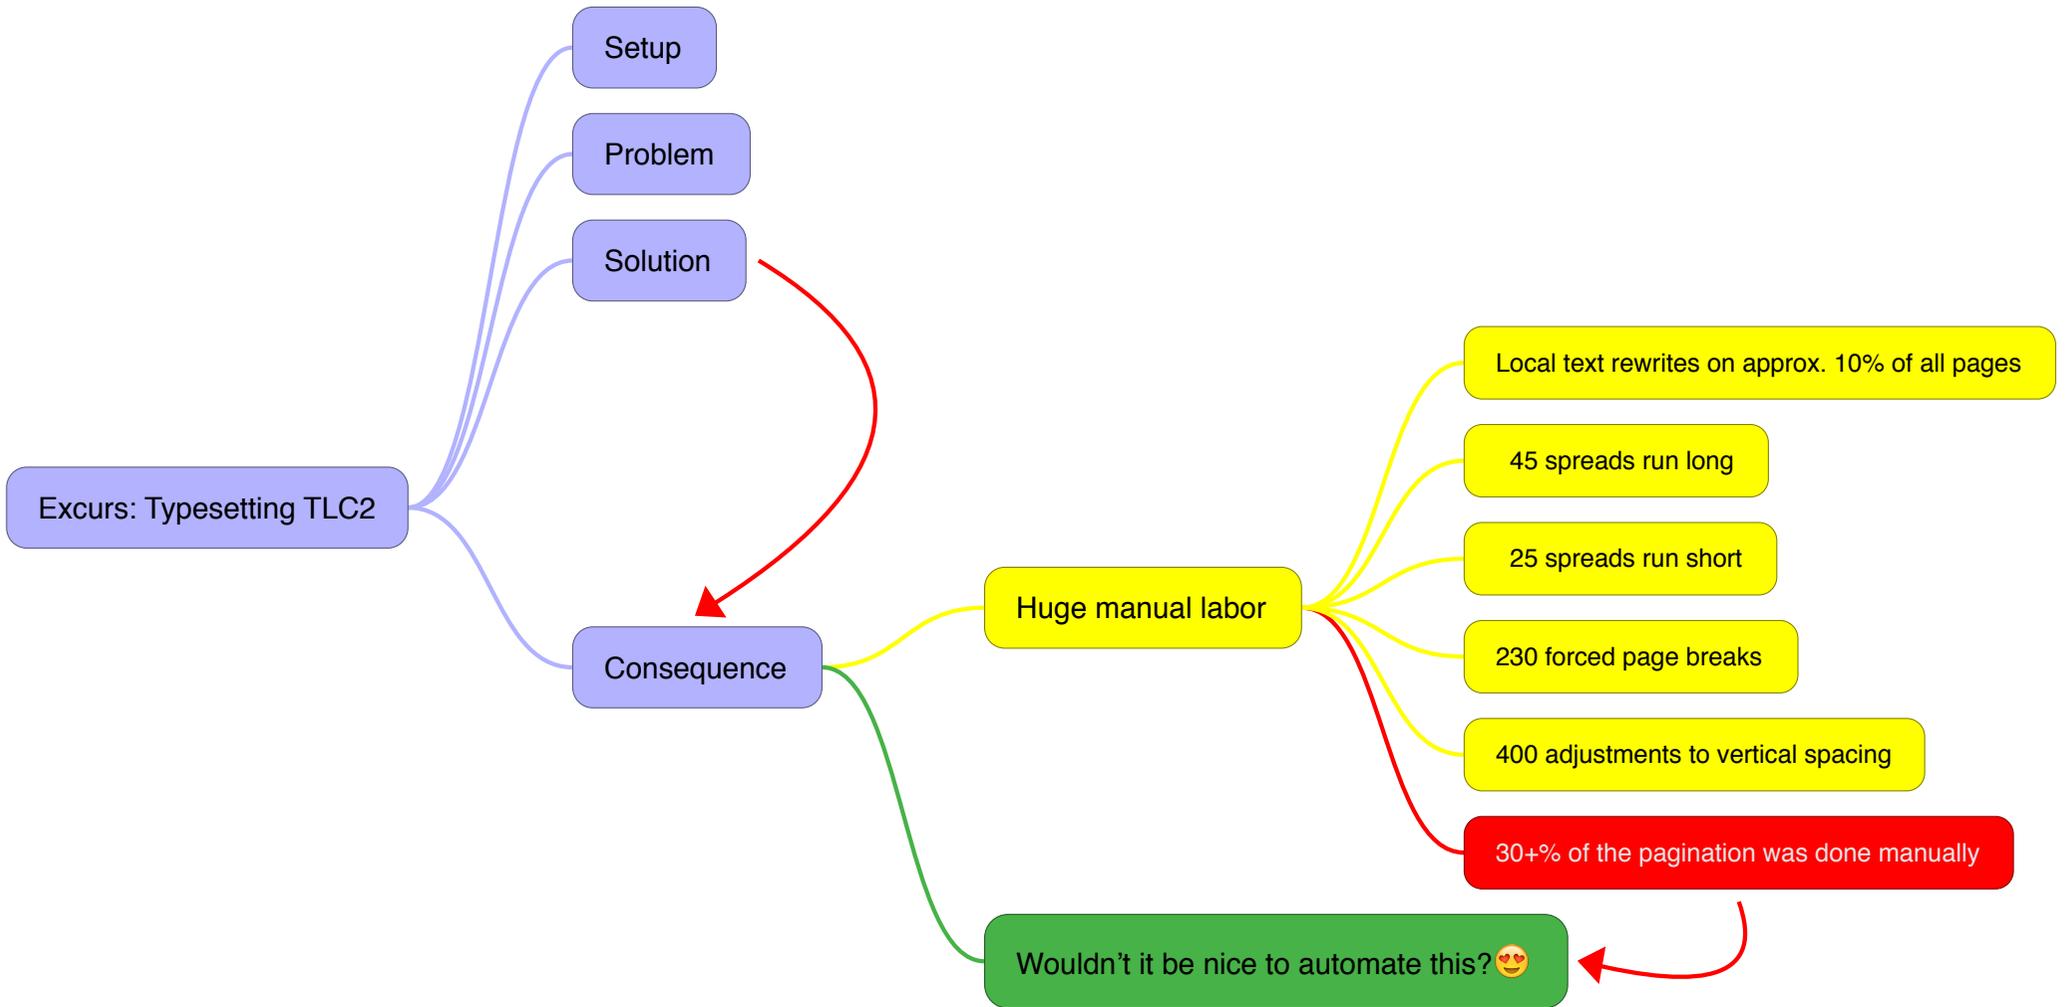
Excurs: Typesetting TLC2













# The pagination problem



## The pagination problem



Today's typesetting engines **all** use greedy pagination algorithms (one page a time)



Why?

The downside of global optimization

Selected research work



Why?

By itself, a **page** (in contrast to a paragraph) **has very limited flexibility** making global optimization fairly pointless in most cases.

Pagination of real life documents has to deal with **unrelated input streams** increasing the complexity drastically

Providing „sensible“ and effective badness measures in case of several input streams is non-trivial and an active research topic

Typesetting requires much more than just pagination: **a usable system also needs to provide the full set of micro-typography** features such as those provided by TeX



Why?

By itself, a **page** (in contrast to a paragraph) **has very limited flexibility** making global optimization fairly pointless in most cases.

Pagination of real life documents has to deal with **unrelated input streams** increasing the complexity drastically

Providing „sensible“ and effective badness measures in case of several input streams is non-trivial and an active research topic

Typesetting requires much more than just pagination: **a usable system also needs to provide the full set of micro-typography** features such as those provided by TeX

unless you make `\baselineskip` flexible which is not really a good option

 Why?

By itself, a **page** (in contrast to a paragraph) **has very limited flexibility** making global optimization fairly pointless in most cases.

Pagination of real life documents has to deal with **unrelated input streams** increasing the complexity drastically

Providing „sensible“ and effective badness measures in case of several input streams is non-trivial and an active research topic

Typesetting requires much more than just pagination: **a usable system also needs to provide the full set of micro-typography** features such as those provided by TeX

Those are: footnotes, marginals, figures, tables, etc.

 Our current solution only considers footnotes (so far) which (obviously) limits the applicability

 A future version of the framework will also address floats!



Why?

By itself, a **page** (in contrast to a paragraph) **has very limited flexibility** making global optimization fairly pointless in most cases.

Pagination of real life documents has to deal with **unrelated input streams** increasing the complexity drastically

Providing „sensible“ and effective badness measures in case of several input streams is non-trivial and an active research topic

Typesetting requires much more than just pagination: **a usable system also needs to provide the full set of micro-typography** features such as those provided by TeX

There are research results, e.g., Plass that show for several measures that they generate NP-hard problems

For example: the badness of a figure placement being measured by the square of the distance to the call-out is NP-hard



Why?

By itself, a **page** (in contrast to a paragraph) **has very limited flexibility** making global optimization fairly pointless in most cases.

Pagination of real life documents has to deal with **unrelated input streams** increasing the complexity drastically

Providing „sensible“ and effective badness measures in case of several input streams is non-trivial and an active research topic

Typesetting requires much more than just pagination: **a usable system also needs to provide the full set of micro-typography features** such as those provided by TeX



This is one reason why all attempts discussed in the research literature ended up with only prototypes that never got any traction (or never got released)



Until recently TeX was difficult to split up into components or interface with — LuaTeX has changed that situation

  
The pagination problem

 Today's typesetting engines **all** use greedy pagination algorithms (one page a time)

 Why?

The downside of global optimization

 Global optimization means that changes can have effects anywhere in the document —not only in later parts

 Dependencies of generated text to pagination (e.g., „see figure X on the following page“) can result in „impossible“ documents or at a minimum in multiple cycles to reach a stable result

Selected research work

  
The pagination problem

 Today's typesetting engines **all** use greedy pagination algorithms (one page a time)

 Why?

The downside of global optimization

Selected research work

1981 Plass: **Optimal Pagination Techniques for Automatic Typesetting Systems**; PhD thesis

1988 Asher: **Type & Set: TeX as the engine of a friendly publishing system**

1998 Wohlfeil: **On the Pagination of Complex Book-Like Documents**; PhD thesis

2003 Jacobs, Li, and Salesin: **Adaptive Document Layout via Manifold Content**

2012 Ciancarini, Furini, and Vitali: **High-quality pagination for publishing**

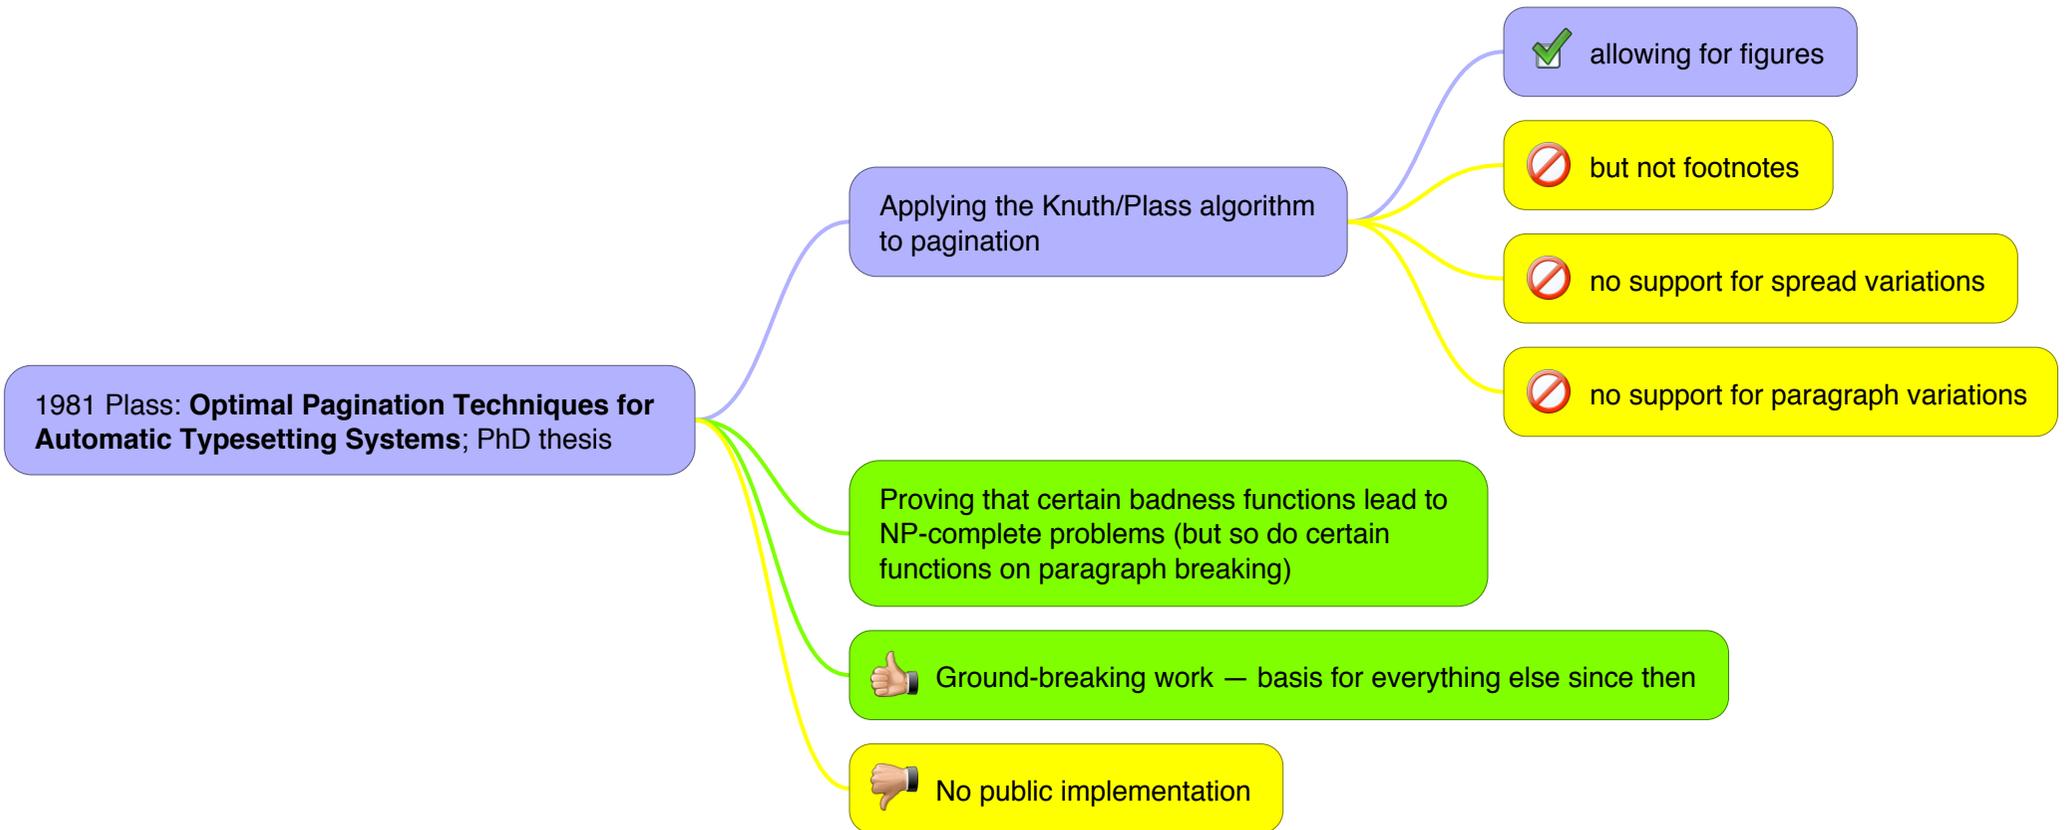
1981 Plass: **Optimal Pagination Techniques for Automatic Typesetting Systems**; PhD thesis

Applying the Knuth/Plass algorithm to pagination

Proving that certain badness functions lead to NP-complete problems (but so do certain functions on paragraph breaking)

 Ground-breaking work — basis for everything else since then

 No public implementation



1988 Asher: **Type & Set: TeX as the engine of a friendly publishing system**



External optimizing pagination on material extracted from `dvi` information  
— exact algorithm used not known



In all likelihood restricted functionality applicable only to a restricted set of documents



No detailed information on the algorithm; no public implementation

1988 Asher: **Type & Set: TeX as the engine of a friendly publishing system**



External optimizing pagination on material extracted from `dvi` information  
— exact algorithm used not known



allowing for figures but apparently limited use in optimization



allowing for footnotes



`dvi` does not contain all information about the available flexibility on the page

Thus limited potential to optimize



In all likelihood restricted functionality applicable only to a restricted set of documents



No detailed information on the algorithm; no public implementation



1998 Wohlfeil: **On the Pagination of Complex Book-Like Documents**; PhD thesis

Applying a variation of the Knuth/Plass algorithm to pagination



Development of „natural“ badness measures for float placement, proving that they can be applied in polynomial time

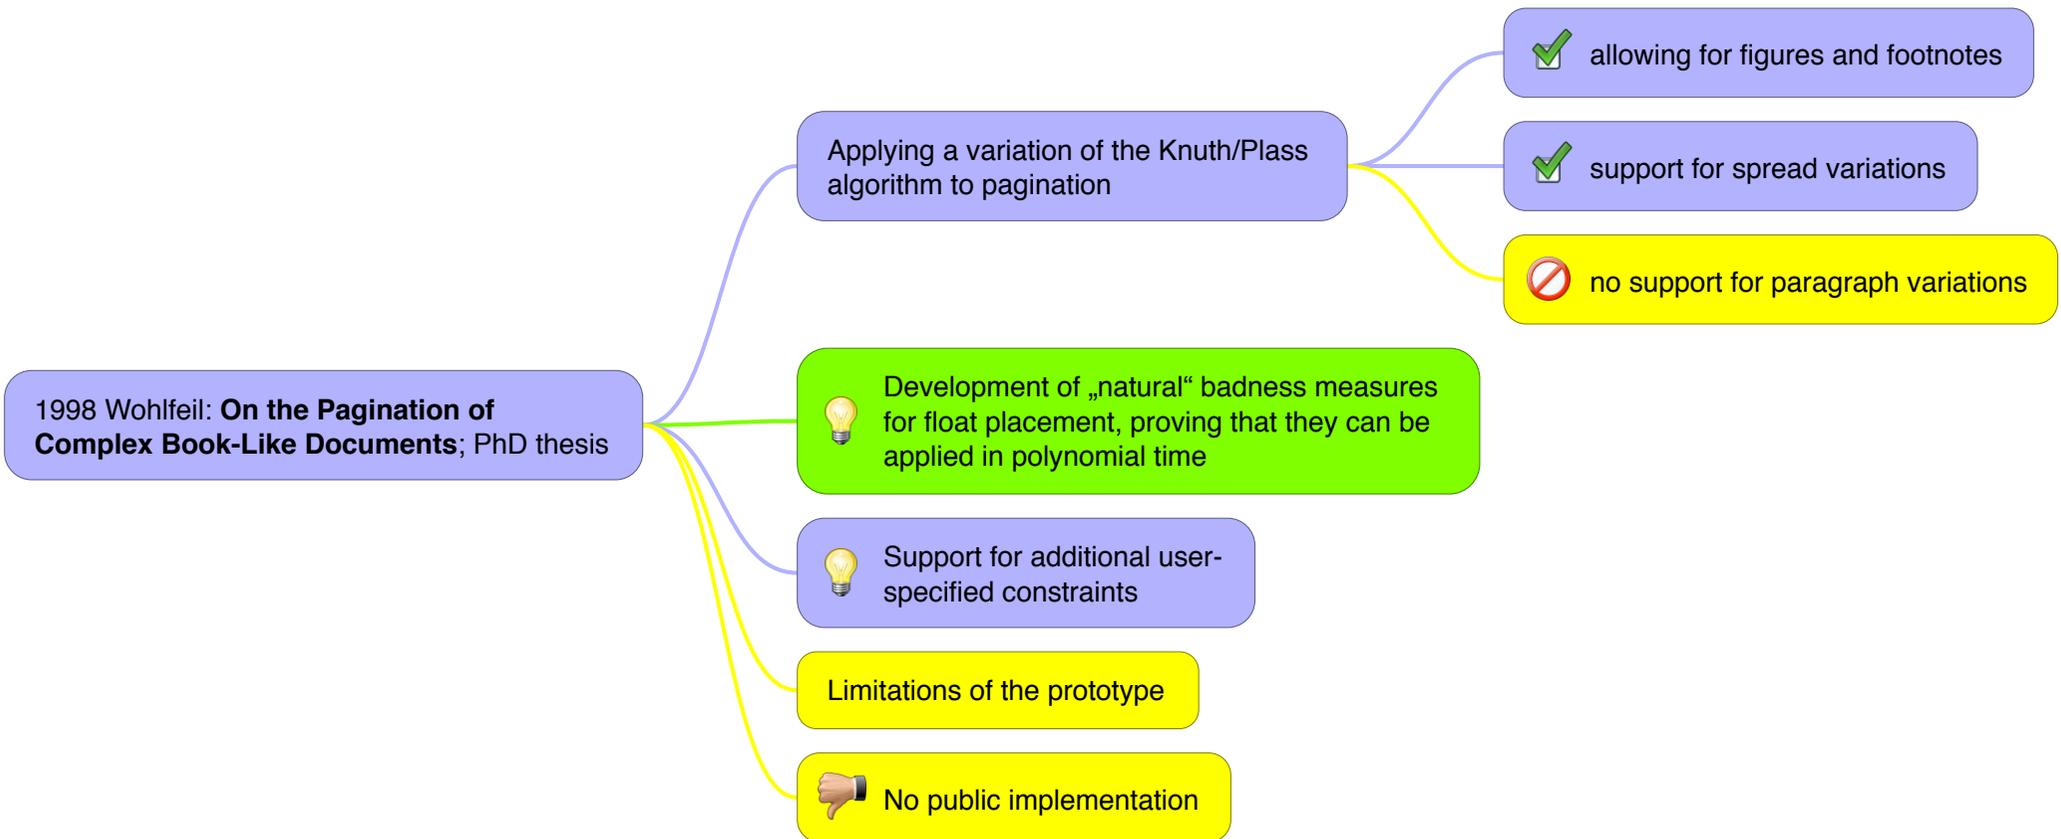


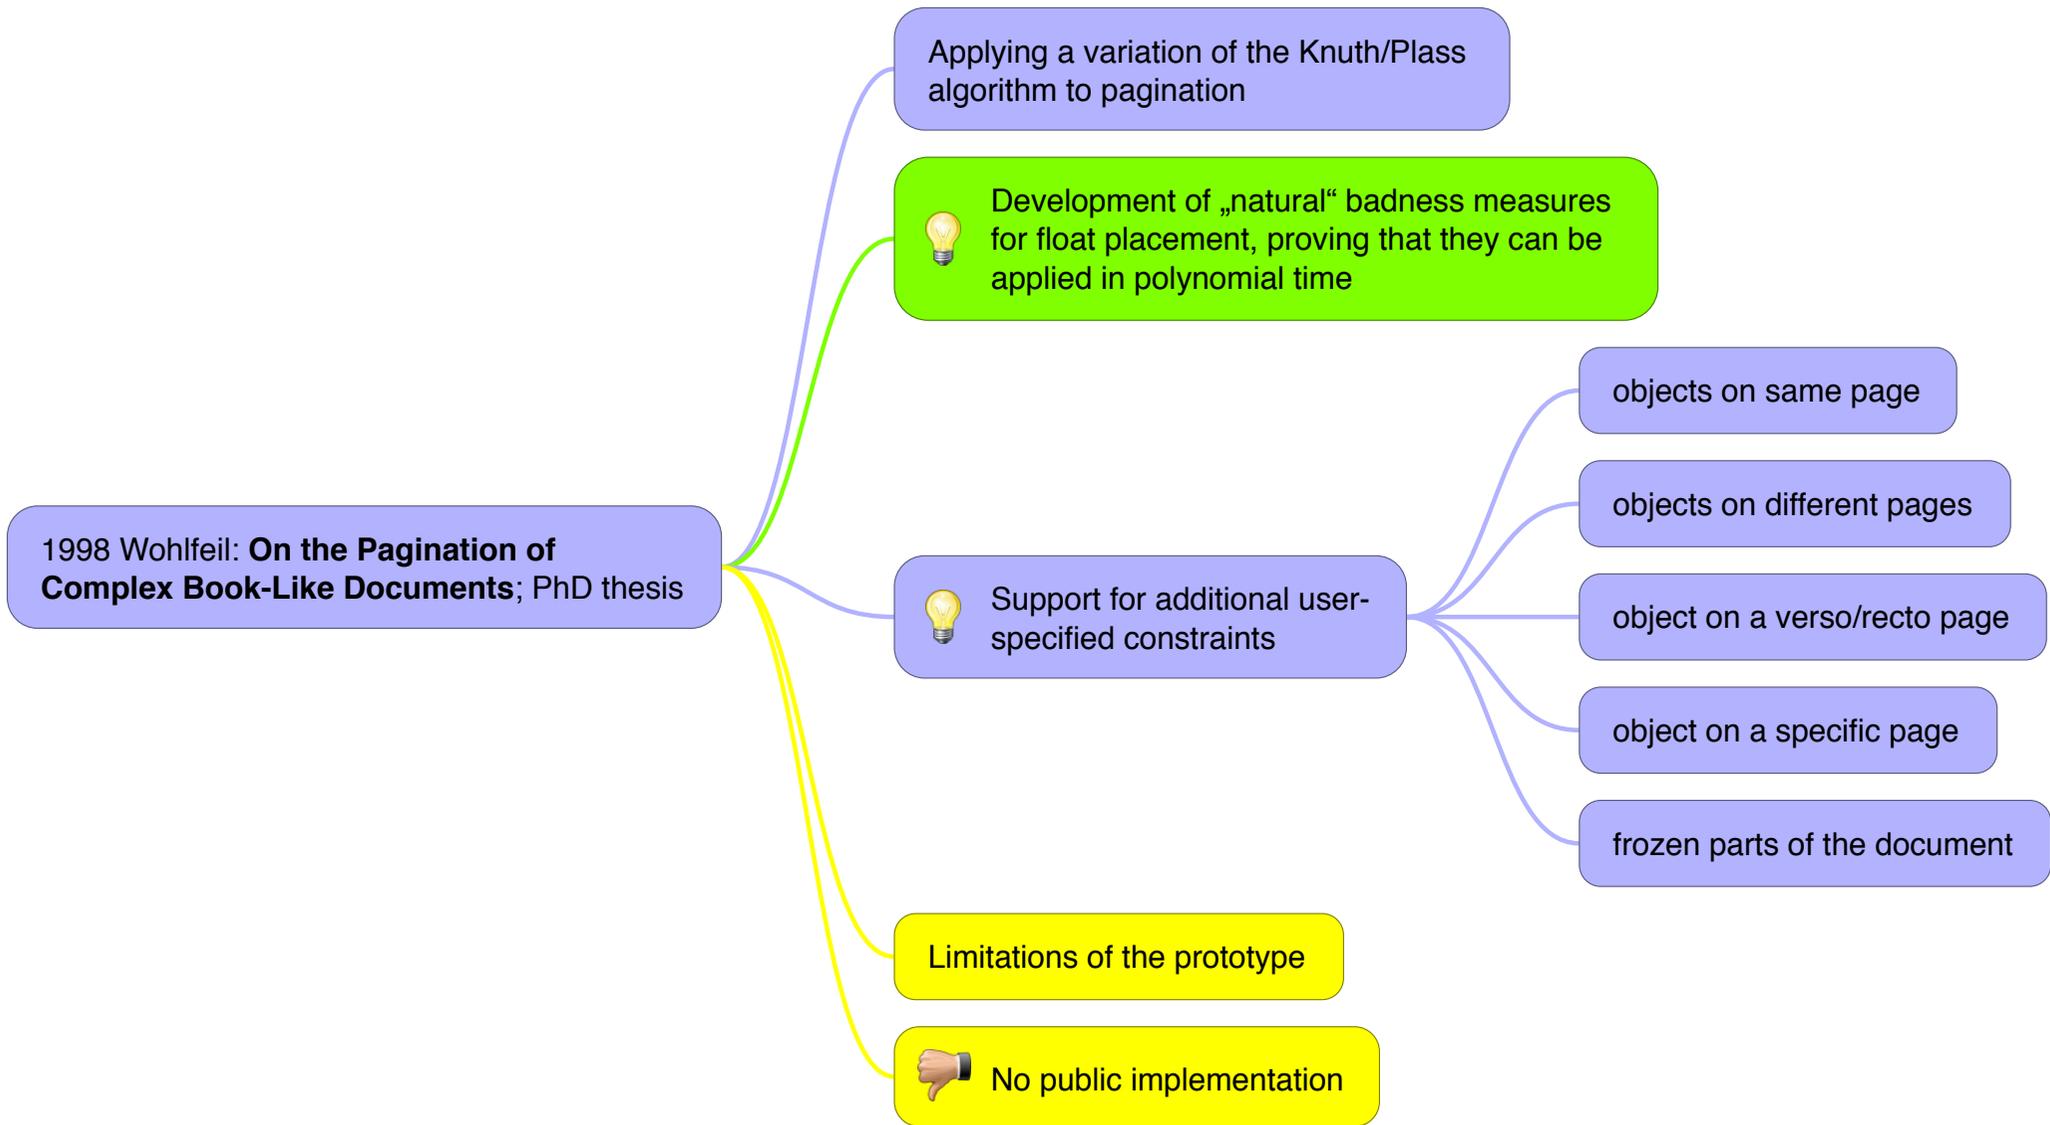
Support for additional user-specified constraints

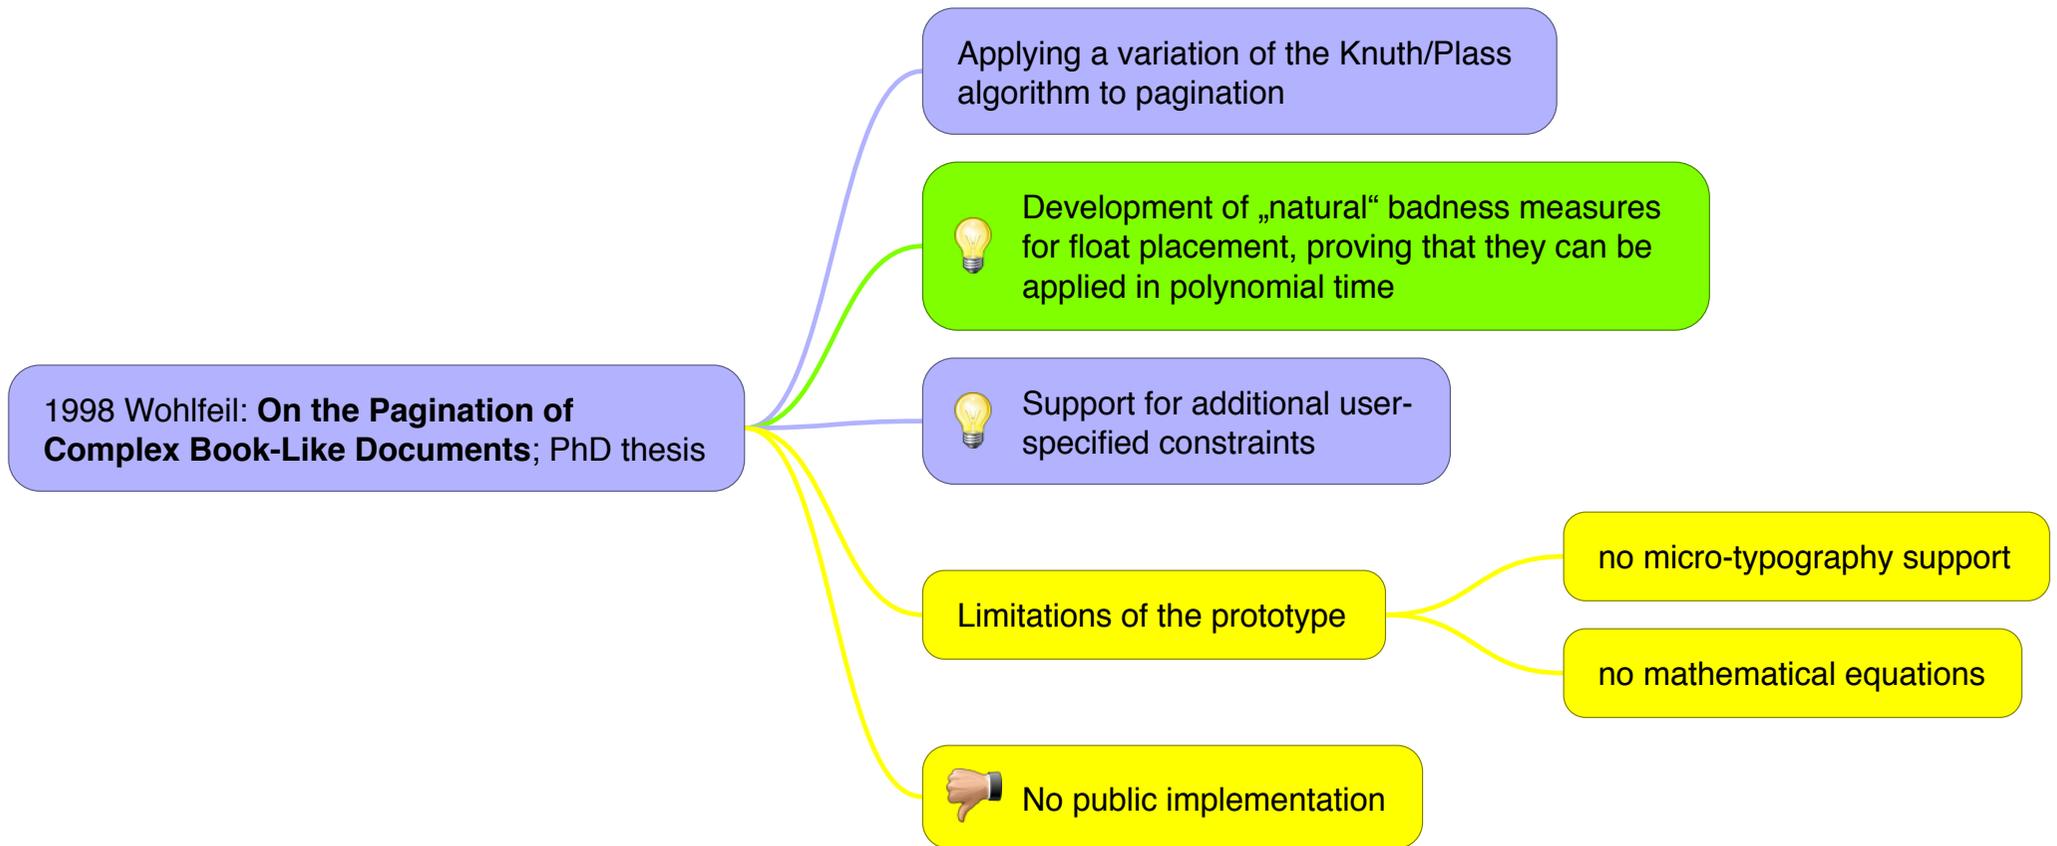
Limitations of the prototype



No public implementation







2003 Jacobs, Li, and Salesin: **Adaptive Document Layout via Manifold Content**



Extension to Knuth/Plass *paragraph* algorithm supporting textual variations (entered through GUI)

Pagination



No public implementation

2003 Jacobs, Li, and Salesin: **Adaptive Document Layout via Manifold Content**

Pagination

💡 Extension to Knuth/Plass *paragraph* algorithm supporting textual variations (entered through GUI)

✅ allowing for figures and footnotes

🚫 no support for spread variations

✅ support for paragraph variations

👎 Essentially greedy: Figure placement is done through brute force trying all combinations

👎 Only local optimization (page by page) using a badness function that incorporates penalties for widows and orphans and non-placed figures

👎 No public implementation

2012 Ciancarini, Furini, and Vitali:  
**High-quality pagination for publishing**

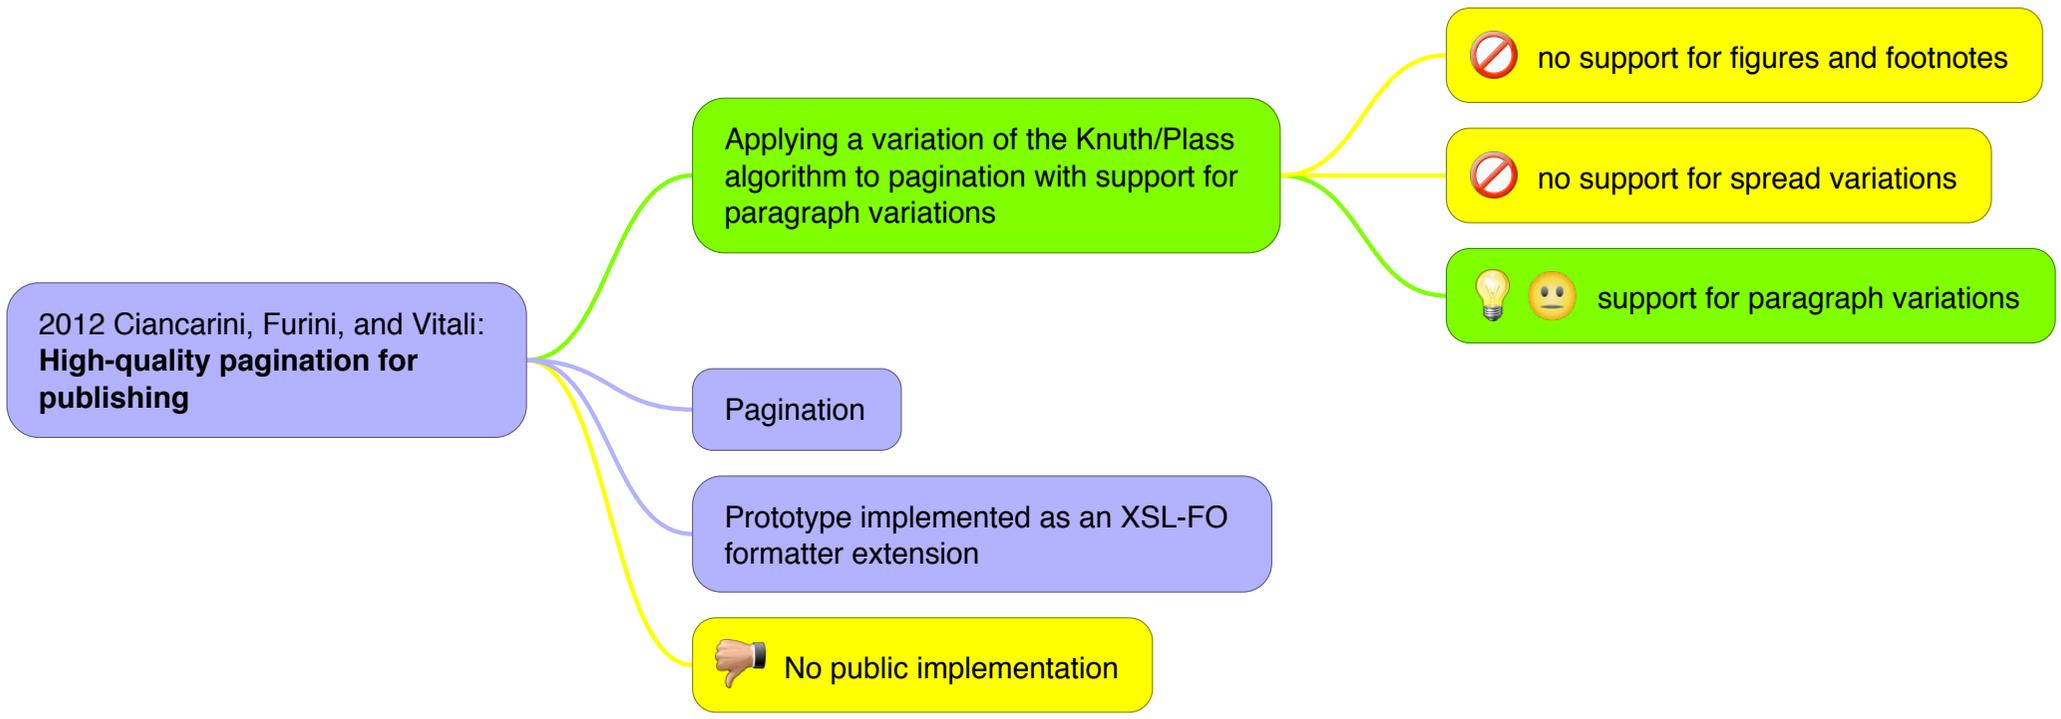
Applying a variation of the Knuth/Plass algorithm to pagination with support for paragraph variations

Pagination

Prototype implemented as an XSL-FO formatter extension



No public implementation



2012 Ciancarini, Furini, and Vitali:  
**High-quality pagination for publishing**

Applying a variation of the Knuth/Plass algorithm to pagination with support for paragraph variations

🚫 no support for figures and footnotes

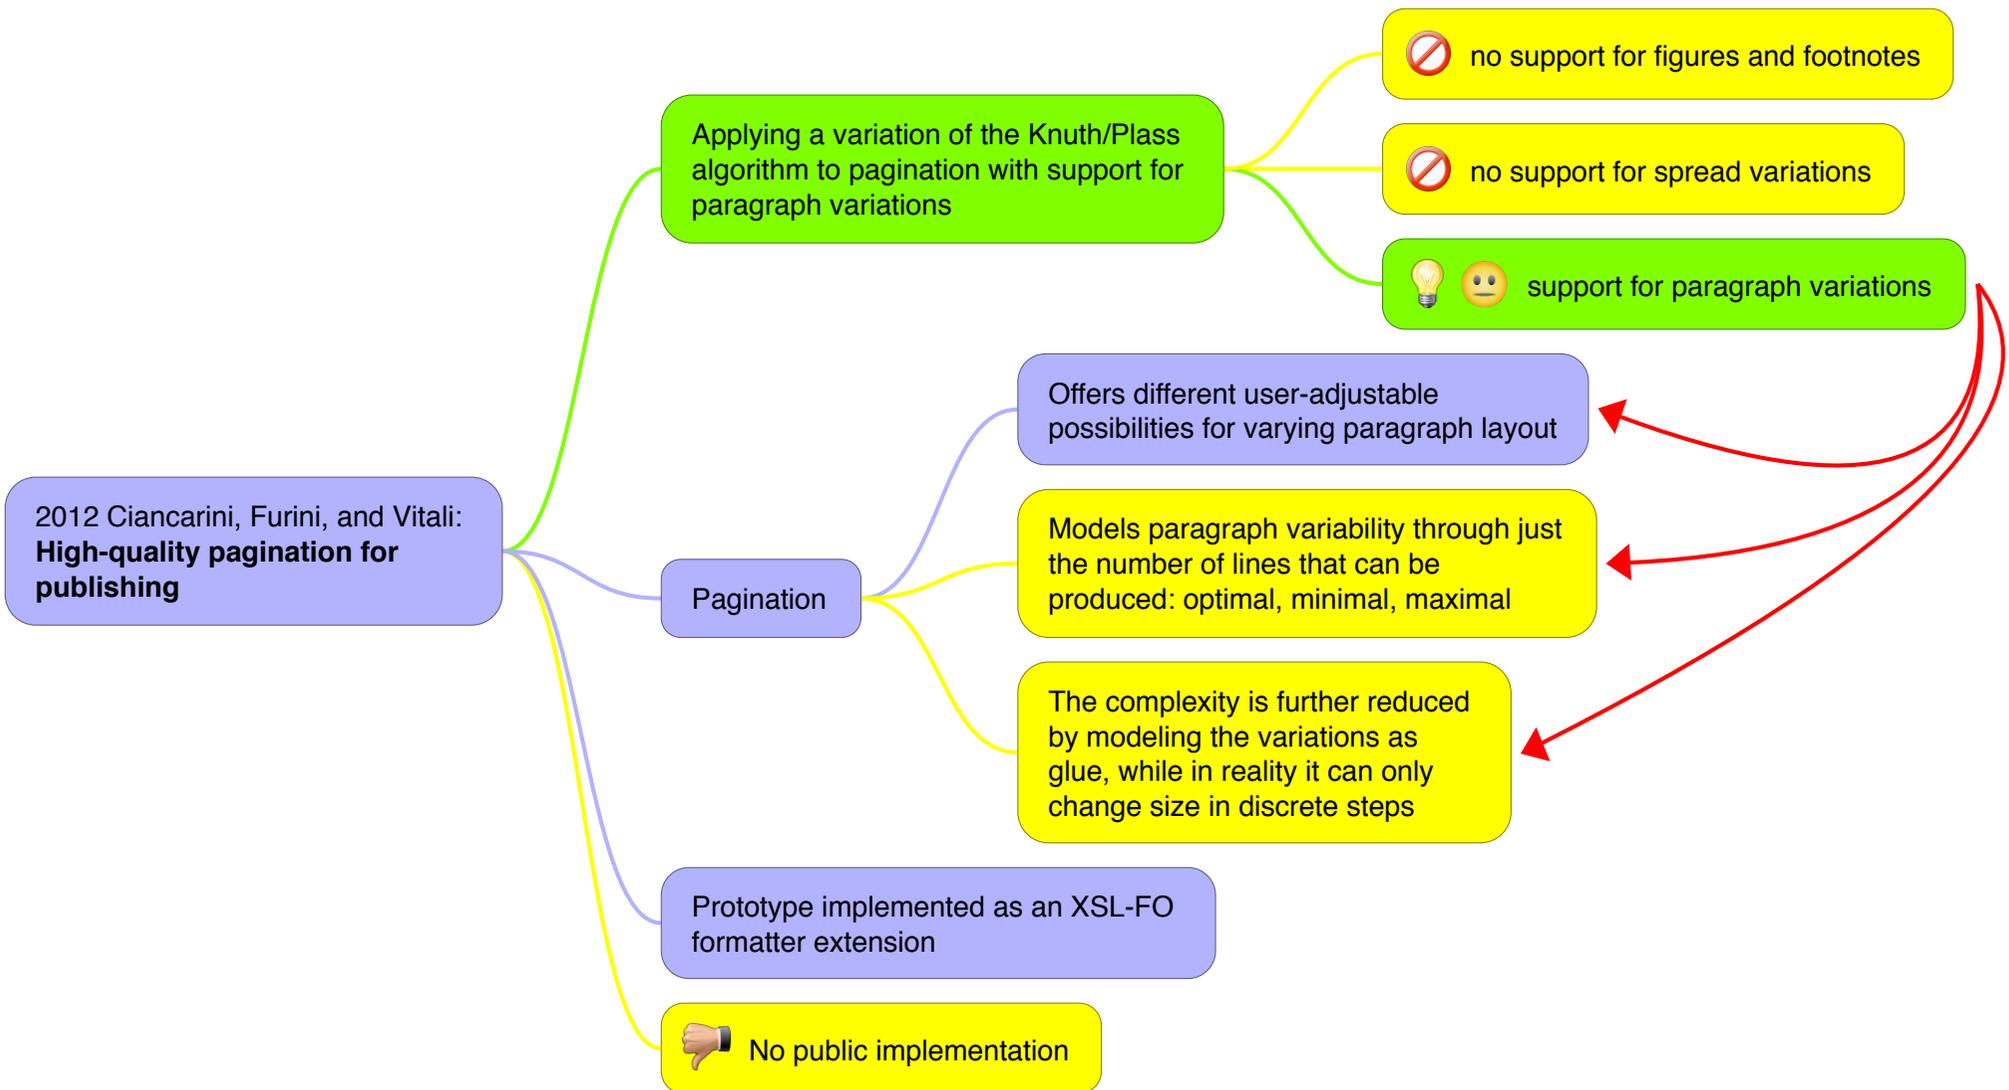
🚫 no support for spread variations

💡 😐 support for paragraph variations

Pagination

Prototype implemented as an XSL-FO formatter extension

👎 No public implementation



2012 Ciancarini, Furini, and Vitali:  
**High-quality pagination for publishing**

Applying a variation of the Knuth/Plass algorithm to pagination with support for paragraph variations

Pagination

Offers different user-adjustable possibilities for varying paragraph layout

Models paragraph variability through just the number of lines that can be produced: optimal, minimal, maximal

The complexity is further reduced by modeling the variations as glue, while in reality it can only change size in discrete steps

☹️ This loses important information as the quality may change drastically between variations or individual lines may get penalized differently for page breaking

Prototype implemented as an XSL-FO formatter extension

👉 No public implementation

2012 Ciancarini, Furini, and Vitali:  
**High-quality pagination for publishing**

Applying a variation of the Knuth/Plass algorithm to pagination with support for paragraph variations

Pagination

Offers different user-adjustable possibilities for varying paragraph layout

Models paragraph variability through just the number of lines that can be produced: optimal, minimal, maximal

The complexity is further reduced by modeling the variations as glue, while in reality it can only change size in discrete steps



This can (and will lead) to very uneven distributions of white space in situations where on a single page all paragraphs are effectively shorter/larger than the expected measure and thus the error is accumulating

Prototype implemented as an XSL-FO formatter extension

No public implementation

2012 Ciancarini, Furini, and Vitali:  
**High-quality pagination for publishing**

Applying a variation of the Knuth/Plass algorithm to pagination with support for paragraph variations

Pagination

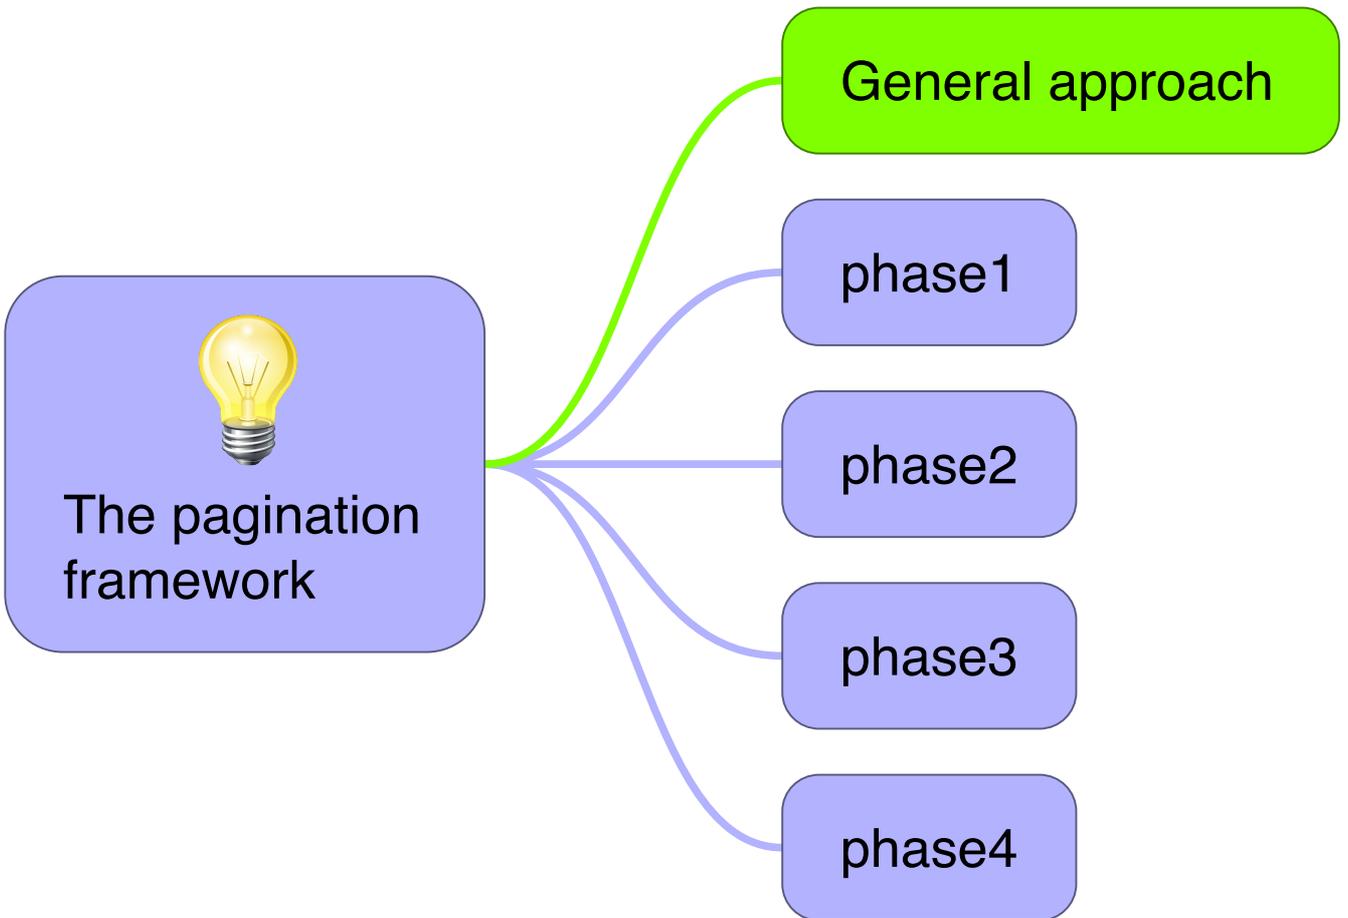
Prototype implemented as an XSL-FO formatter extension

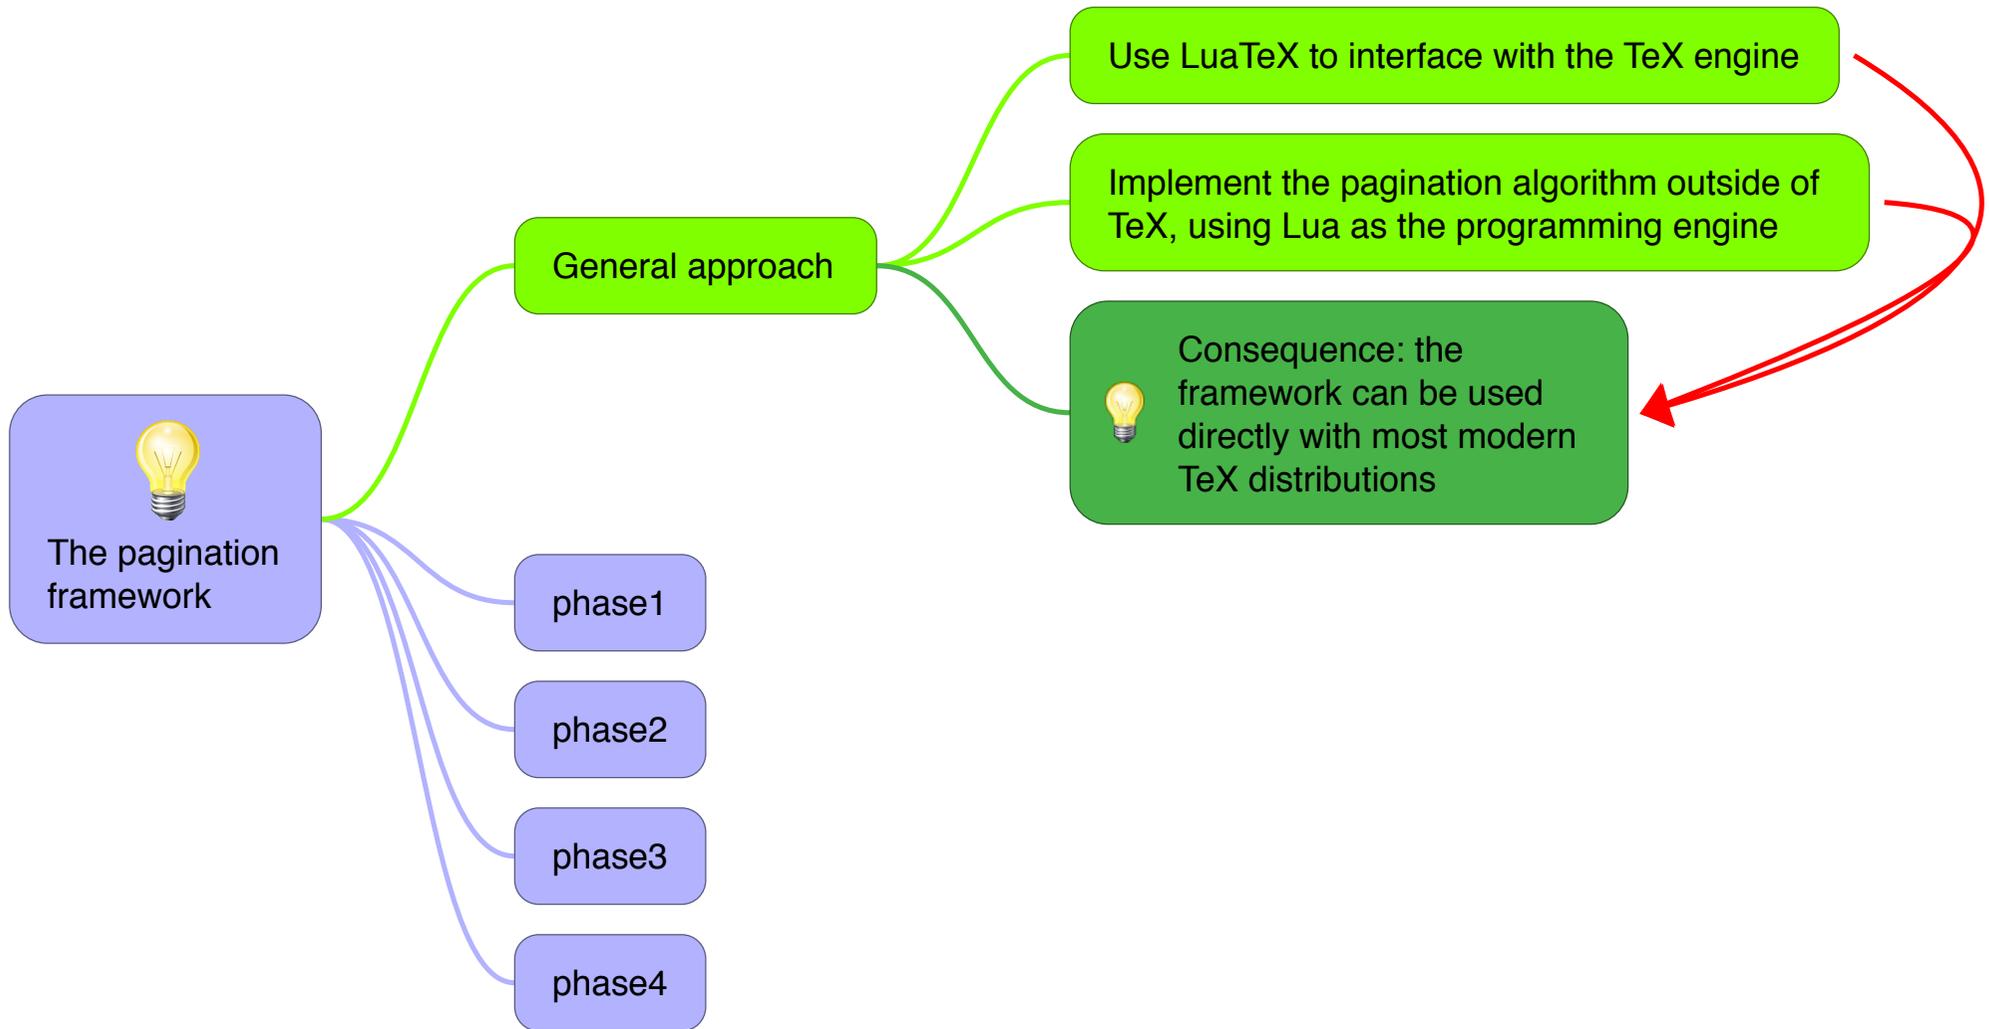


No public implementation



# The pagination framework





  
The pagination framework

phase1

phase2

phase3

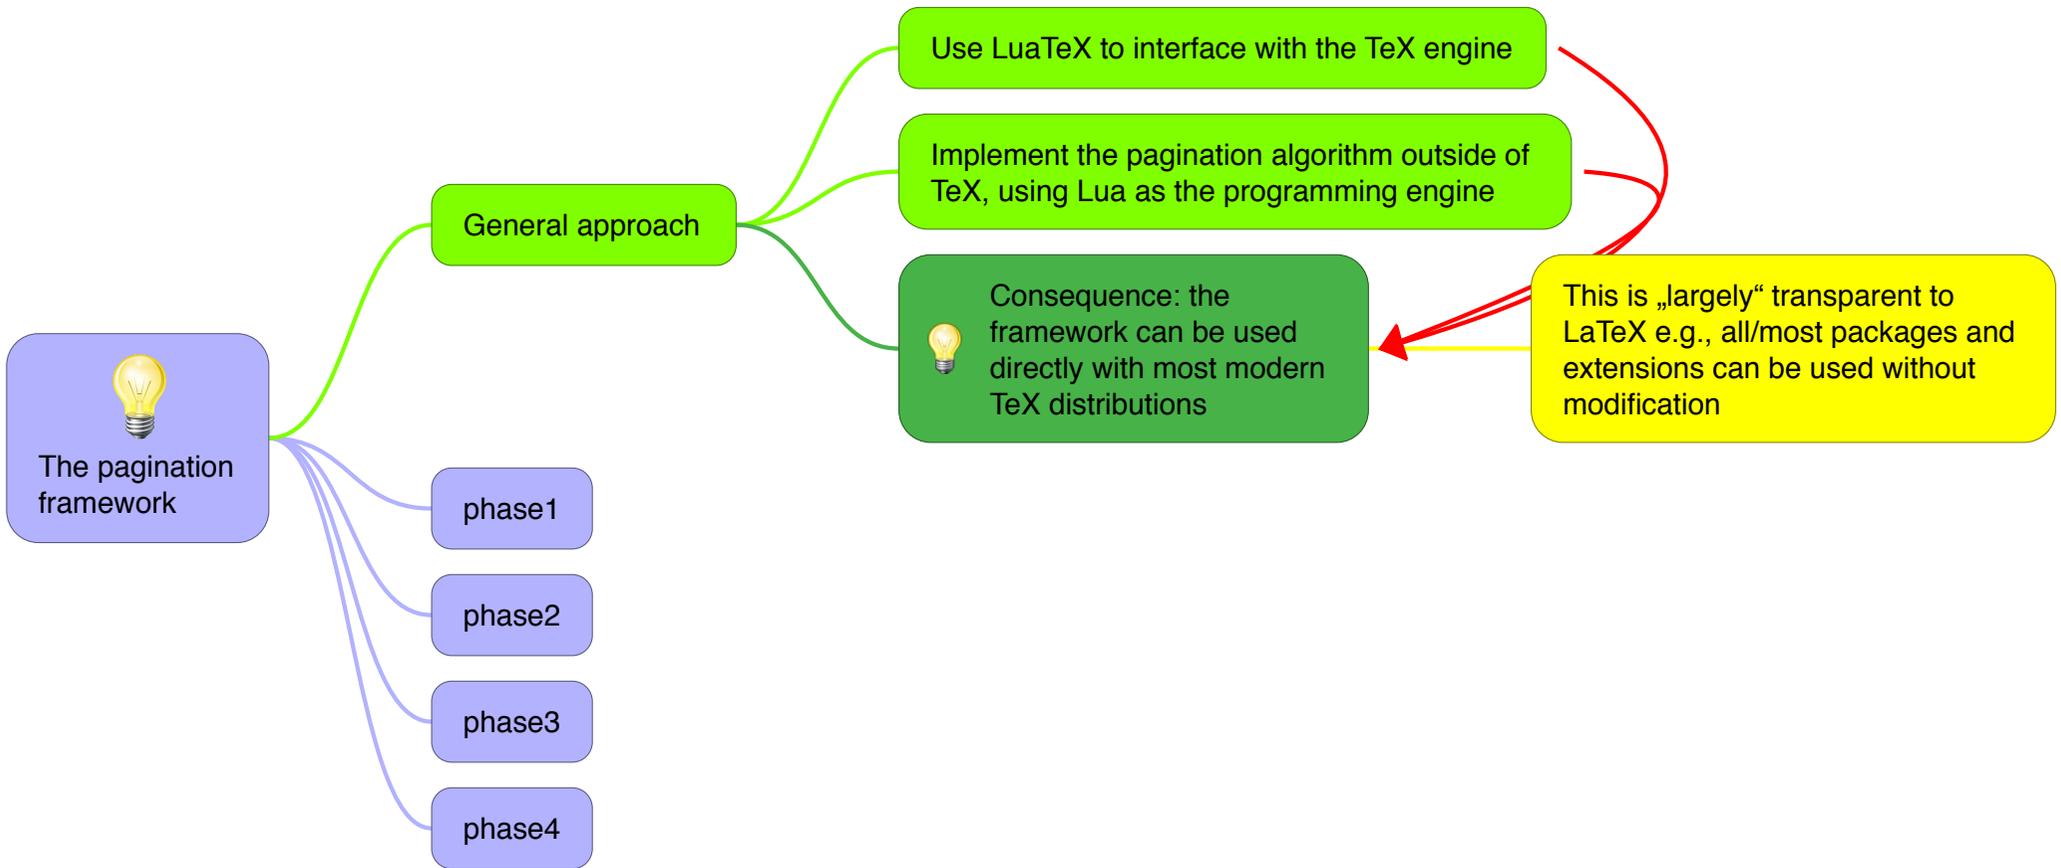
phase4

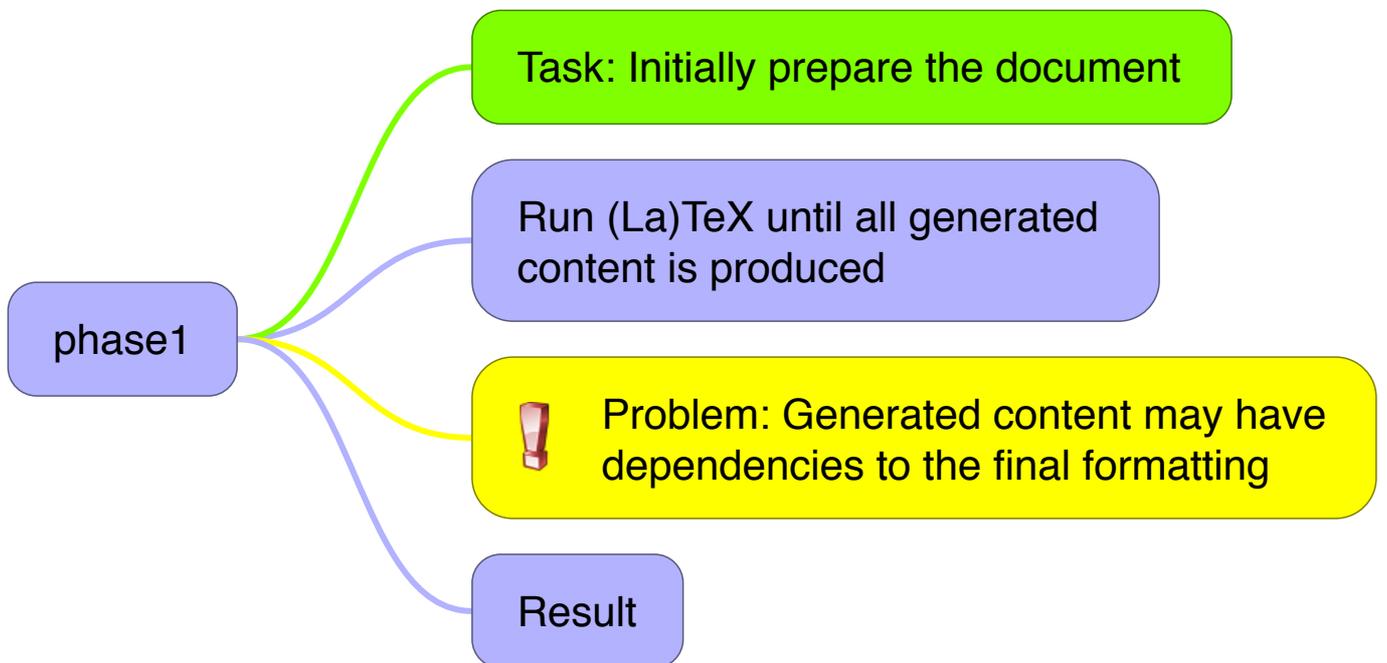
General approach

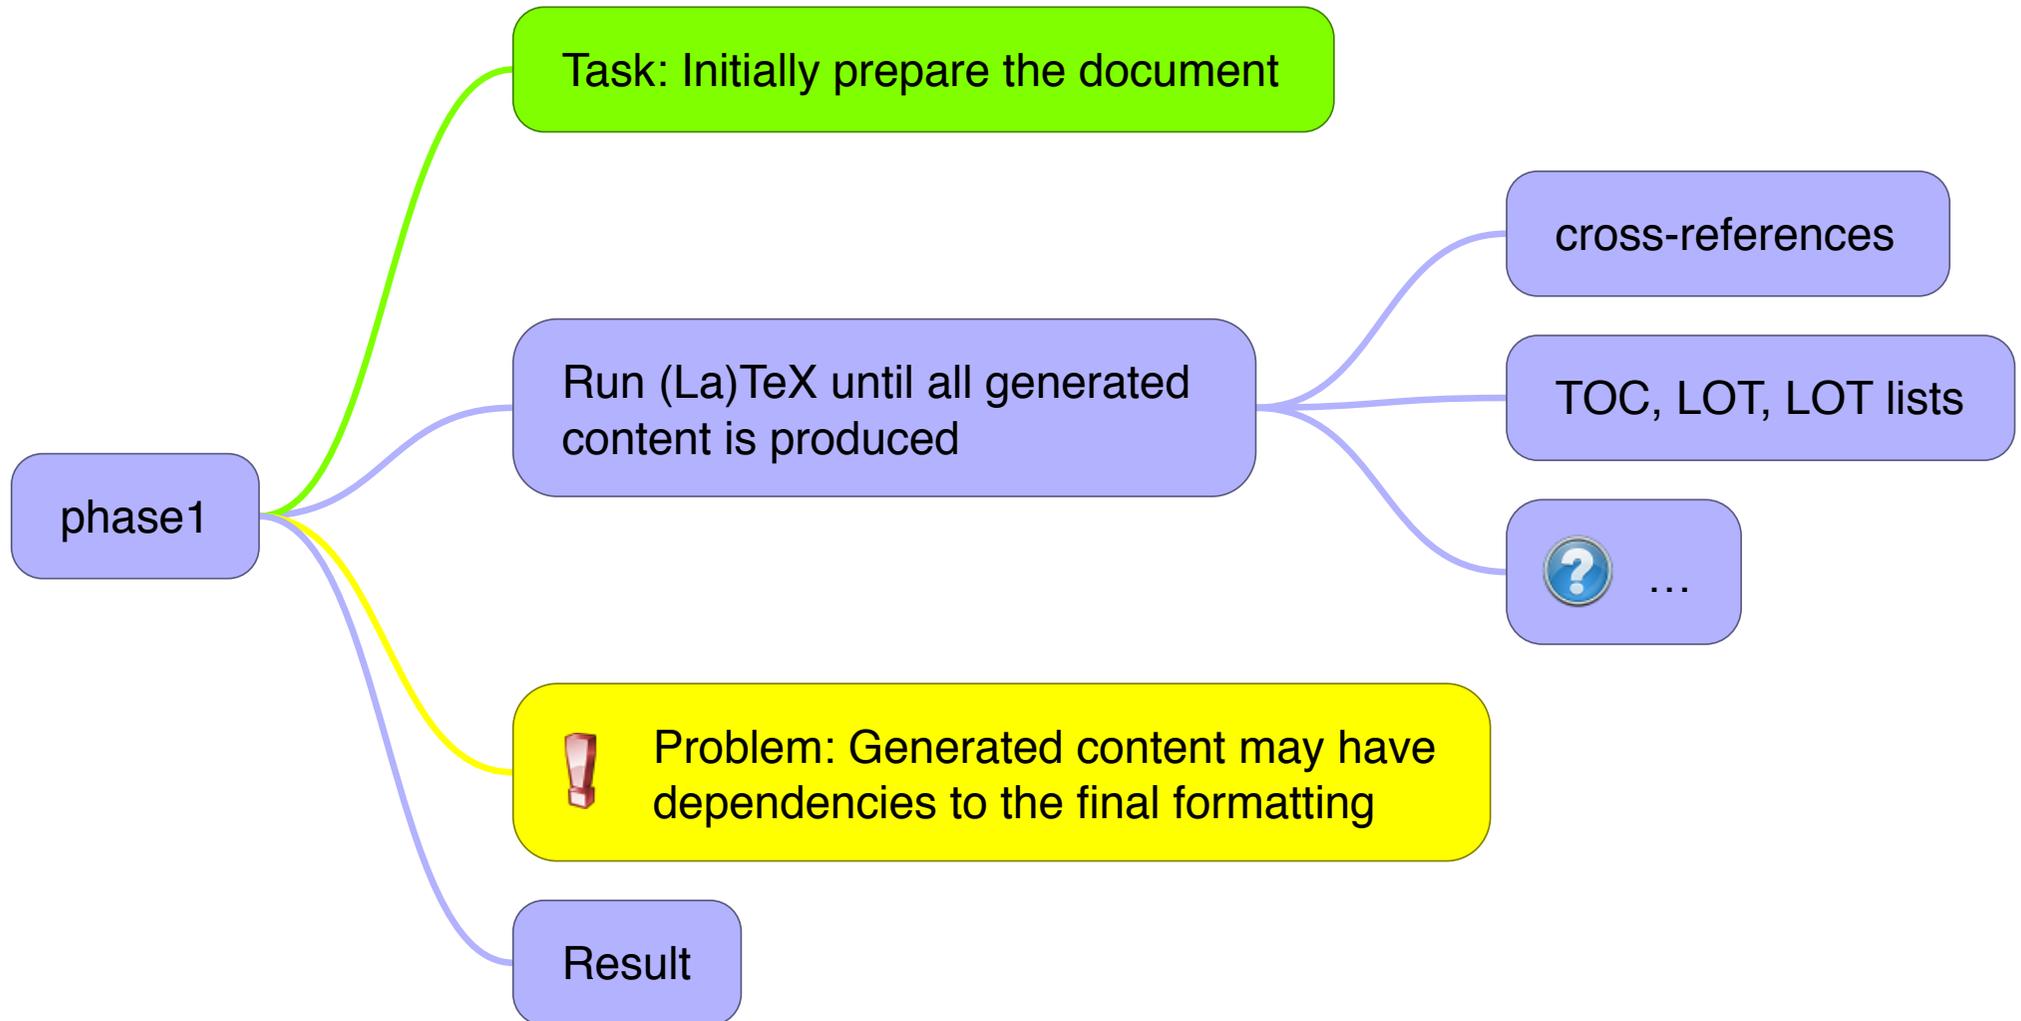
Use LuaTeX to interface with the TeX engine

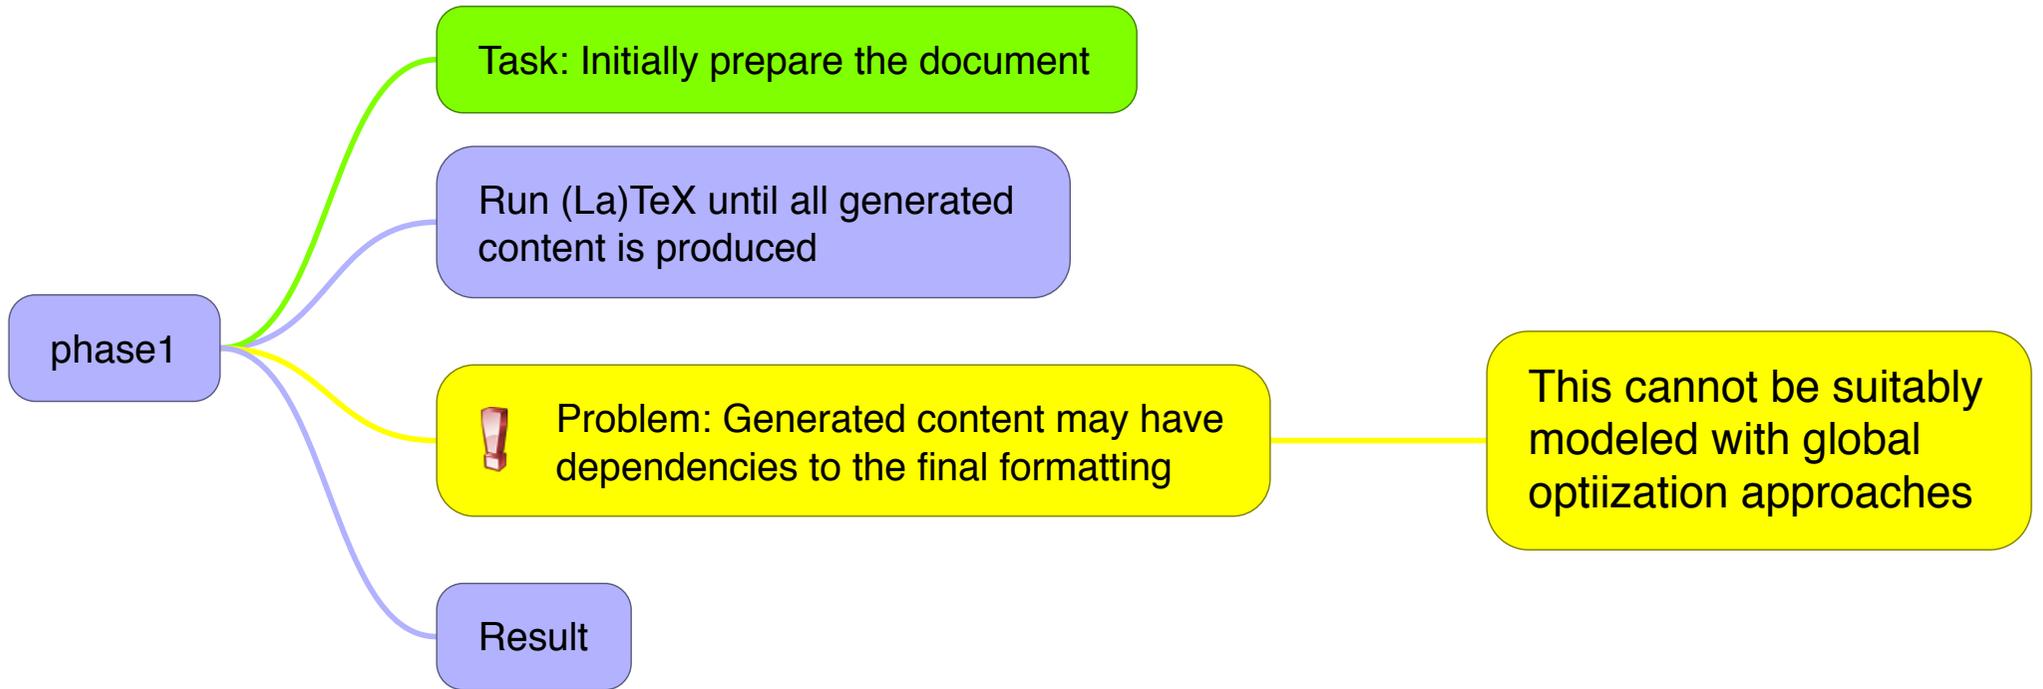
Implement the pagination algorithm outside of TeX, using Lua as the programming engine

 Consequence: the framework can be used directly with most modern TeX distributions









phase1

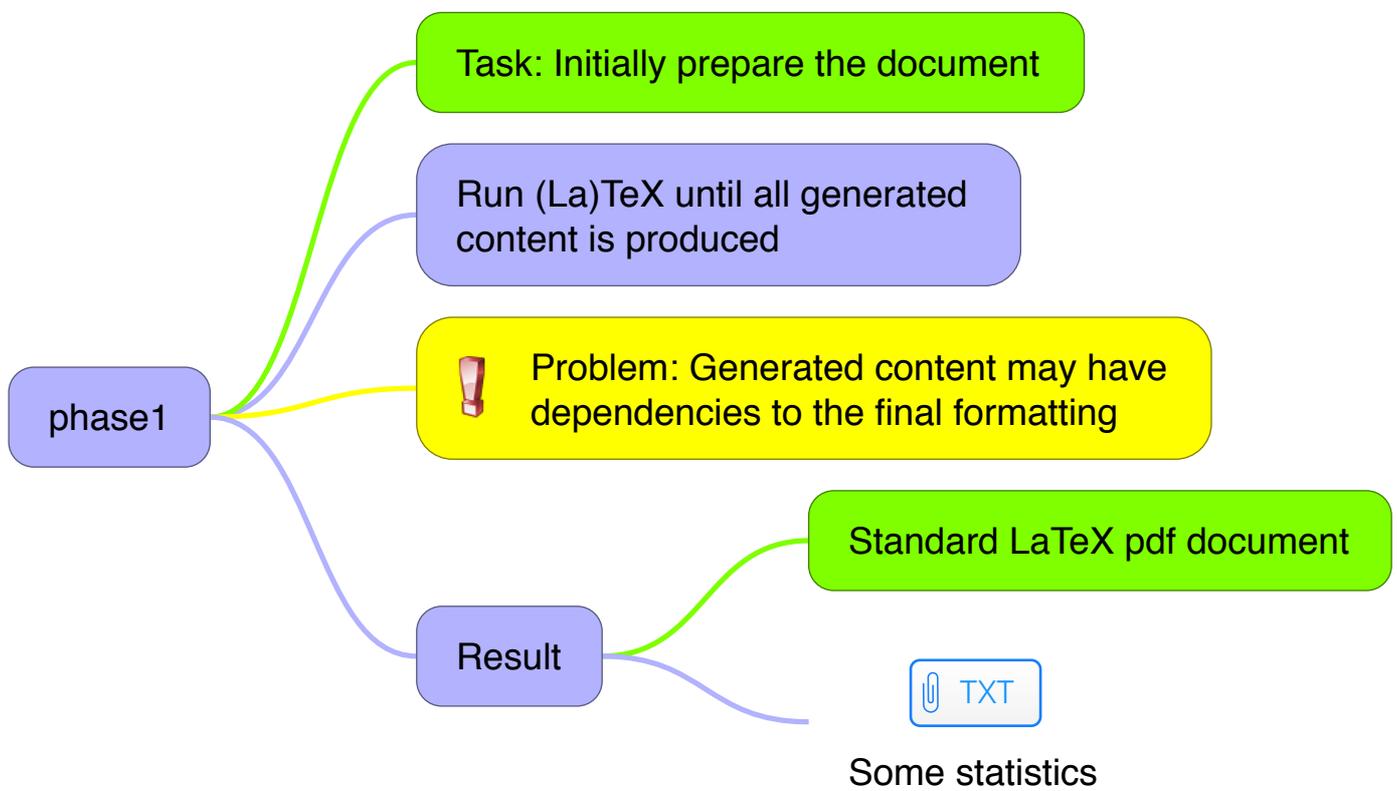
Task: Initially prepare the document

Run (La)TeX until all generated content is produced

! Problem: Generated content may have dependencies to the final formatting

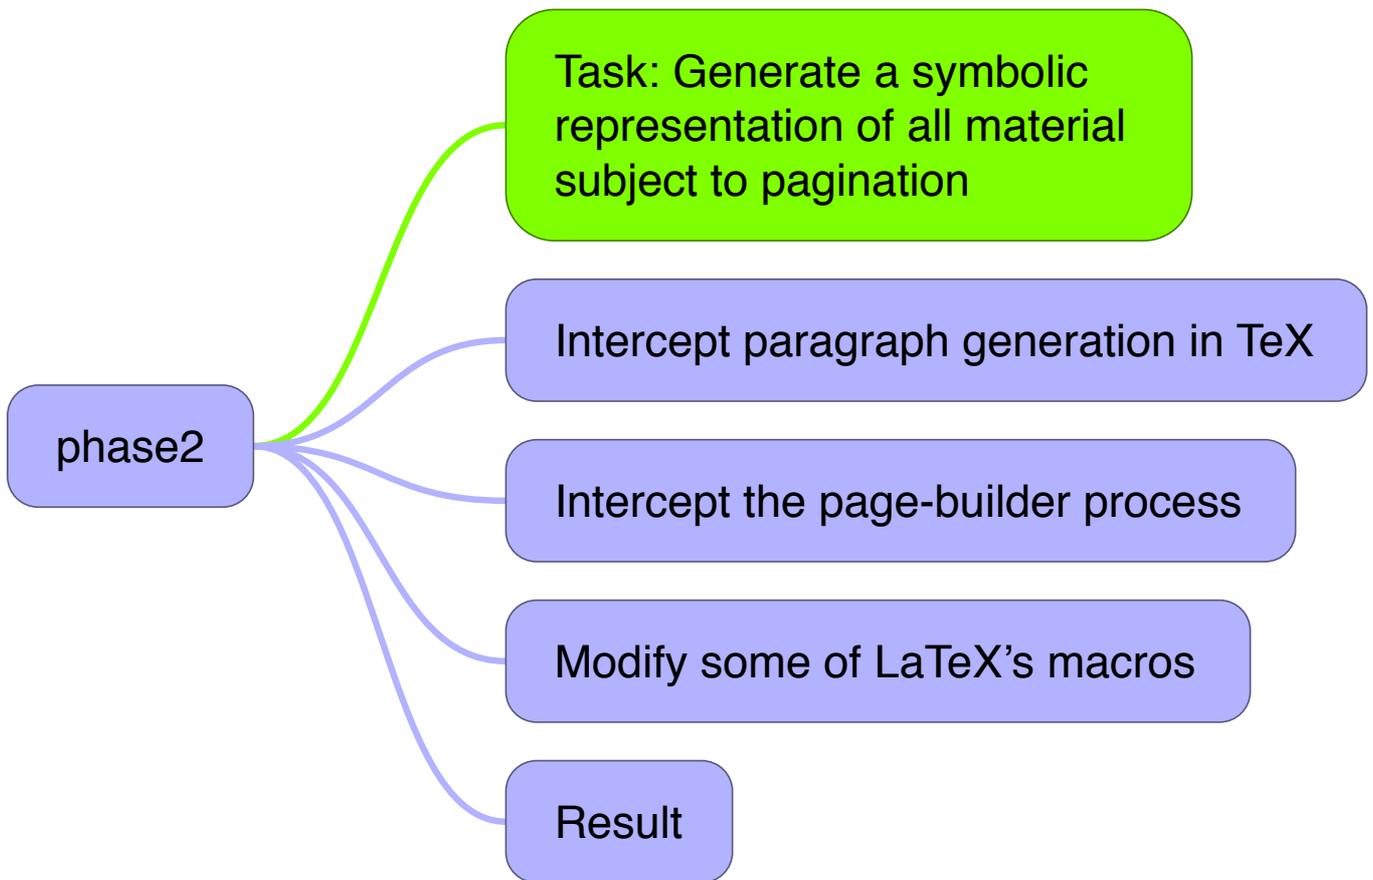
Result

This cannot be suitably modeled with global optimization approaches





**Some statistics**



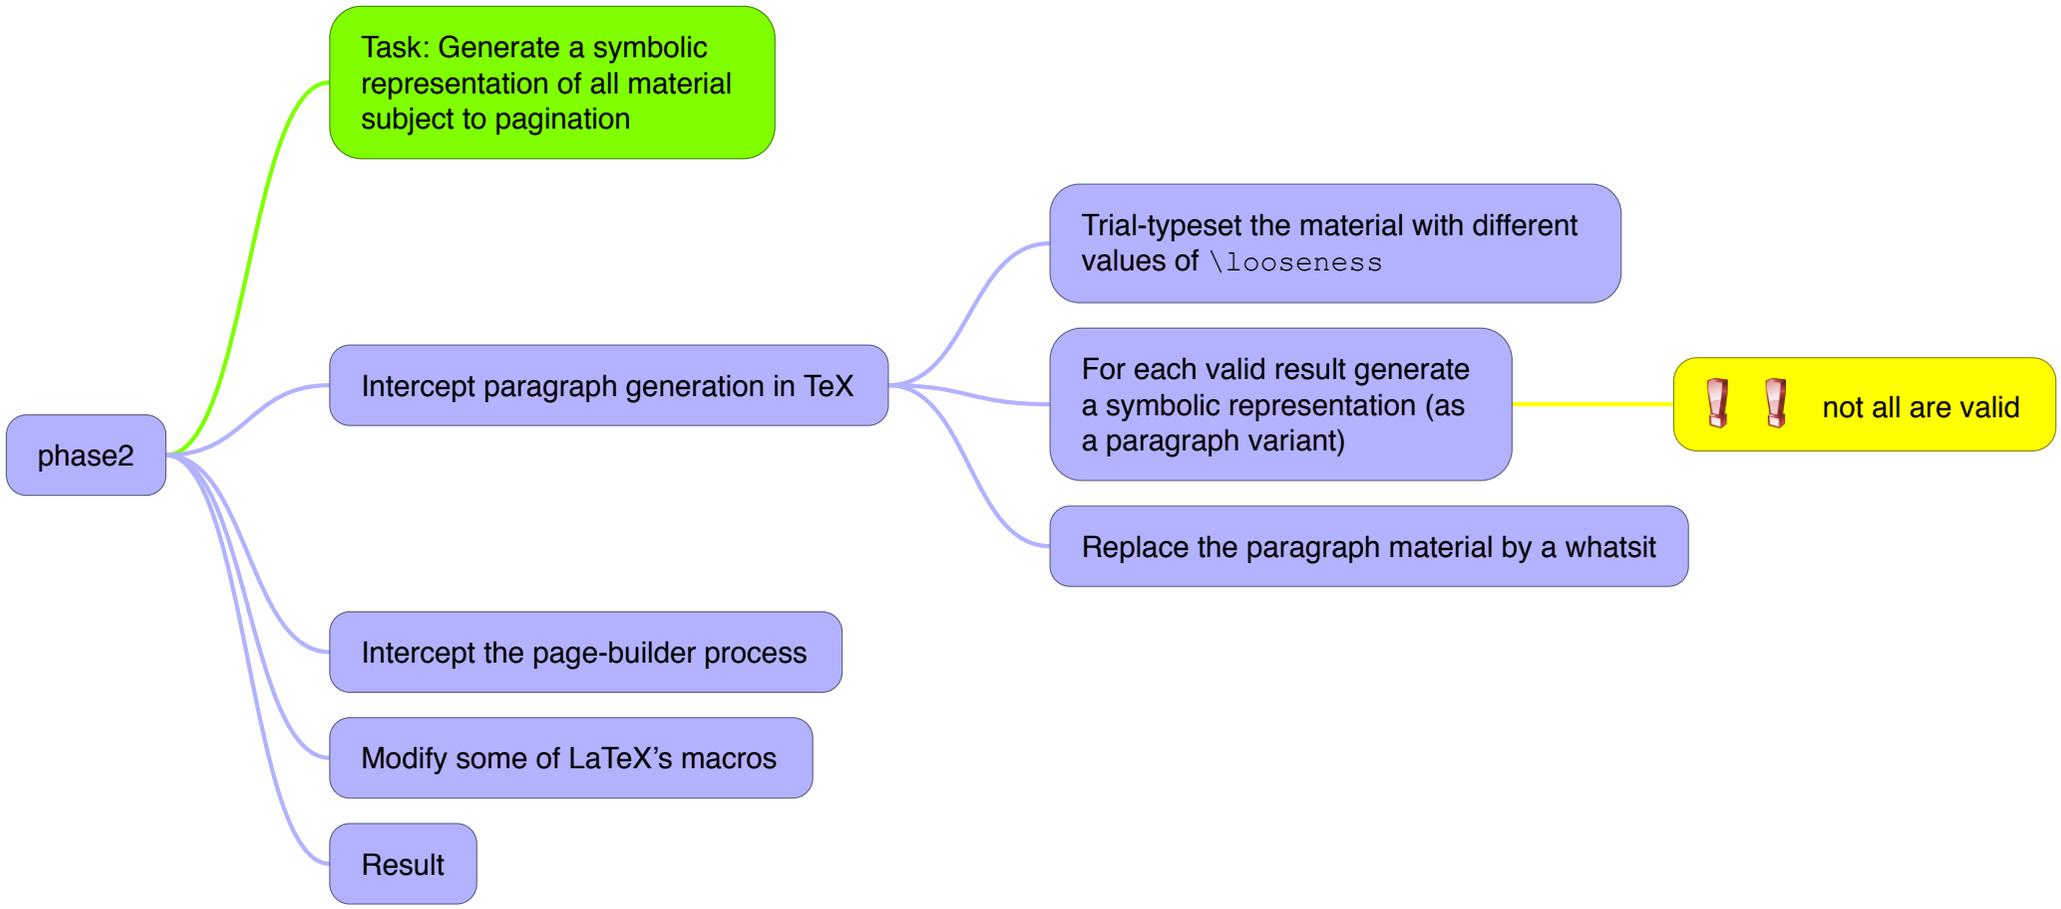
Task: Generate a symbolic representation of all material subject to pagination

`GNode.discardable(<seqno>, <penalty>, <height>, <plus>, <fil>, <fill>, <filll>, <minus>)`

`GNode.normal(<seqno>, <height>, <depth>, <plus>, <fil>, <fill>, <filll>, <minus>)`

`GNode.insert(<seqno>, <height>, <depth>)`

`GNode.special(<seqno>, .., .., ..)`



phase2

Task: Generate a symbolic representation of all material subject to pagination

Intercept paragraph generation in TeX

Intercept the page-builder process

Modify some of LaTeX's macros

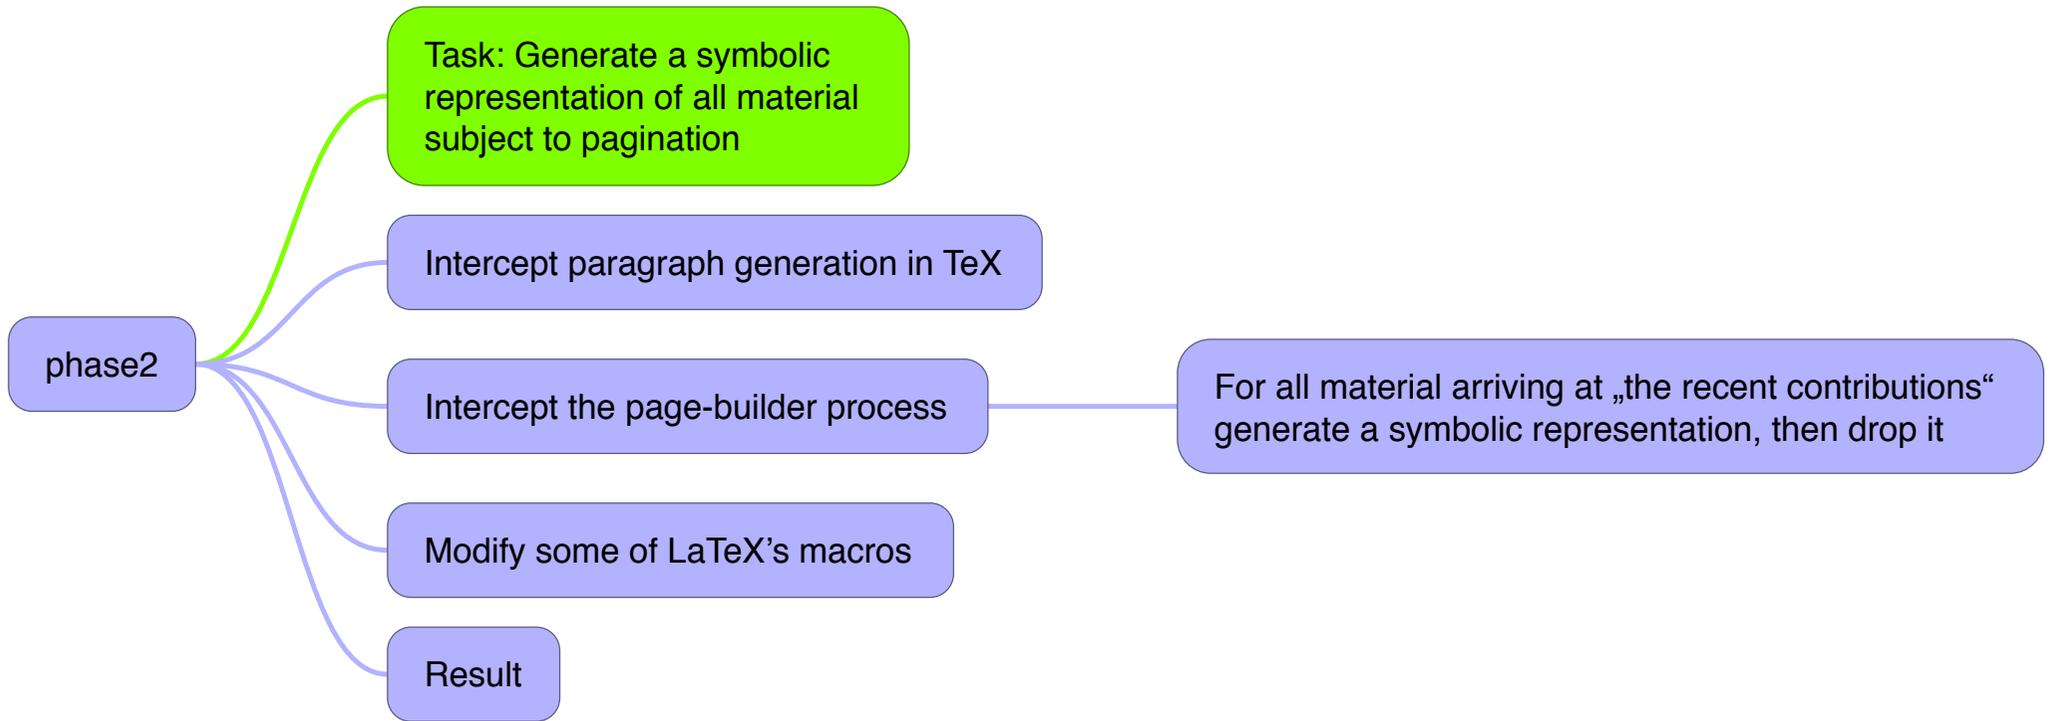
Result

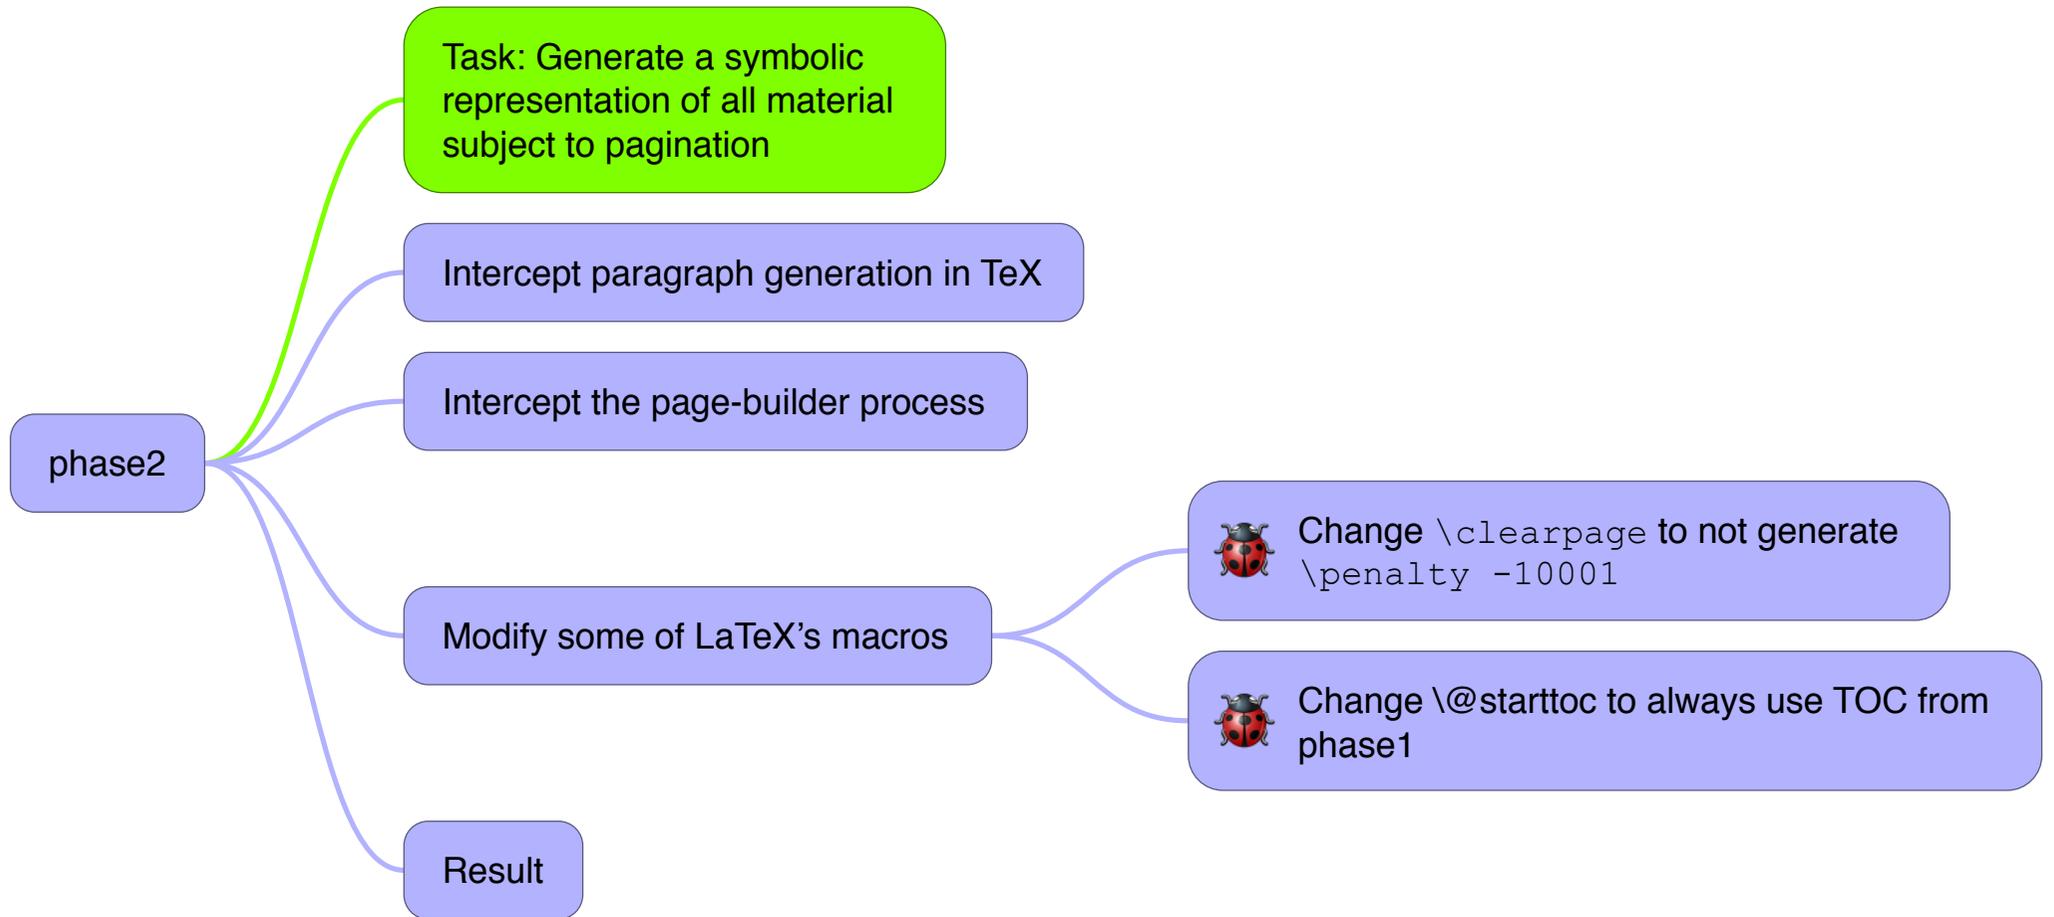
Trial-typeset the material with different values of `\looseness`

For each valid result generate a symbolic representation (as a paragraph variant)

Replace the paragraph material by a whatsit

!! not all are valid





phase2

Task: Generate a symbolic representation of all material subject to pagination

Intercept paragraph generation in TeX

Intercept the page-builder process

Modify some of LaTeX's macros

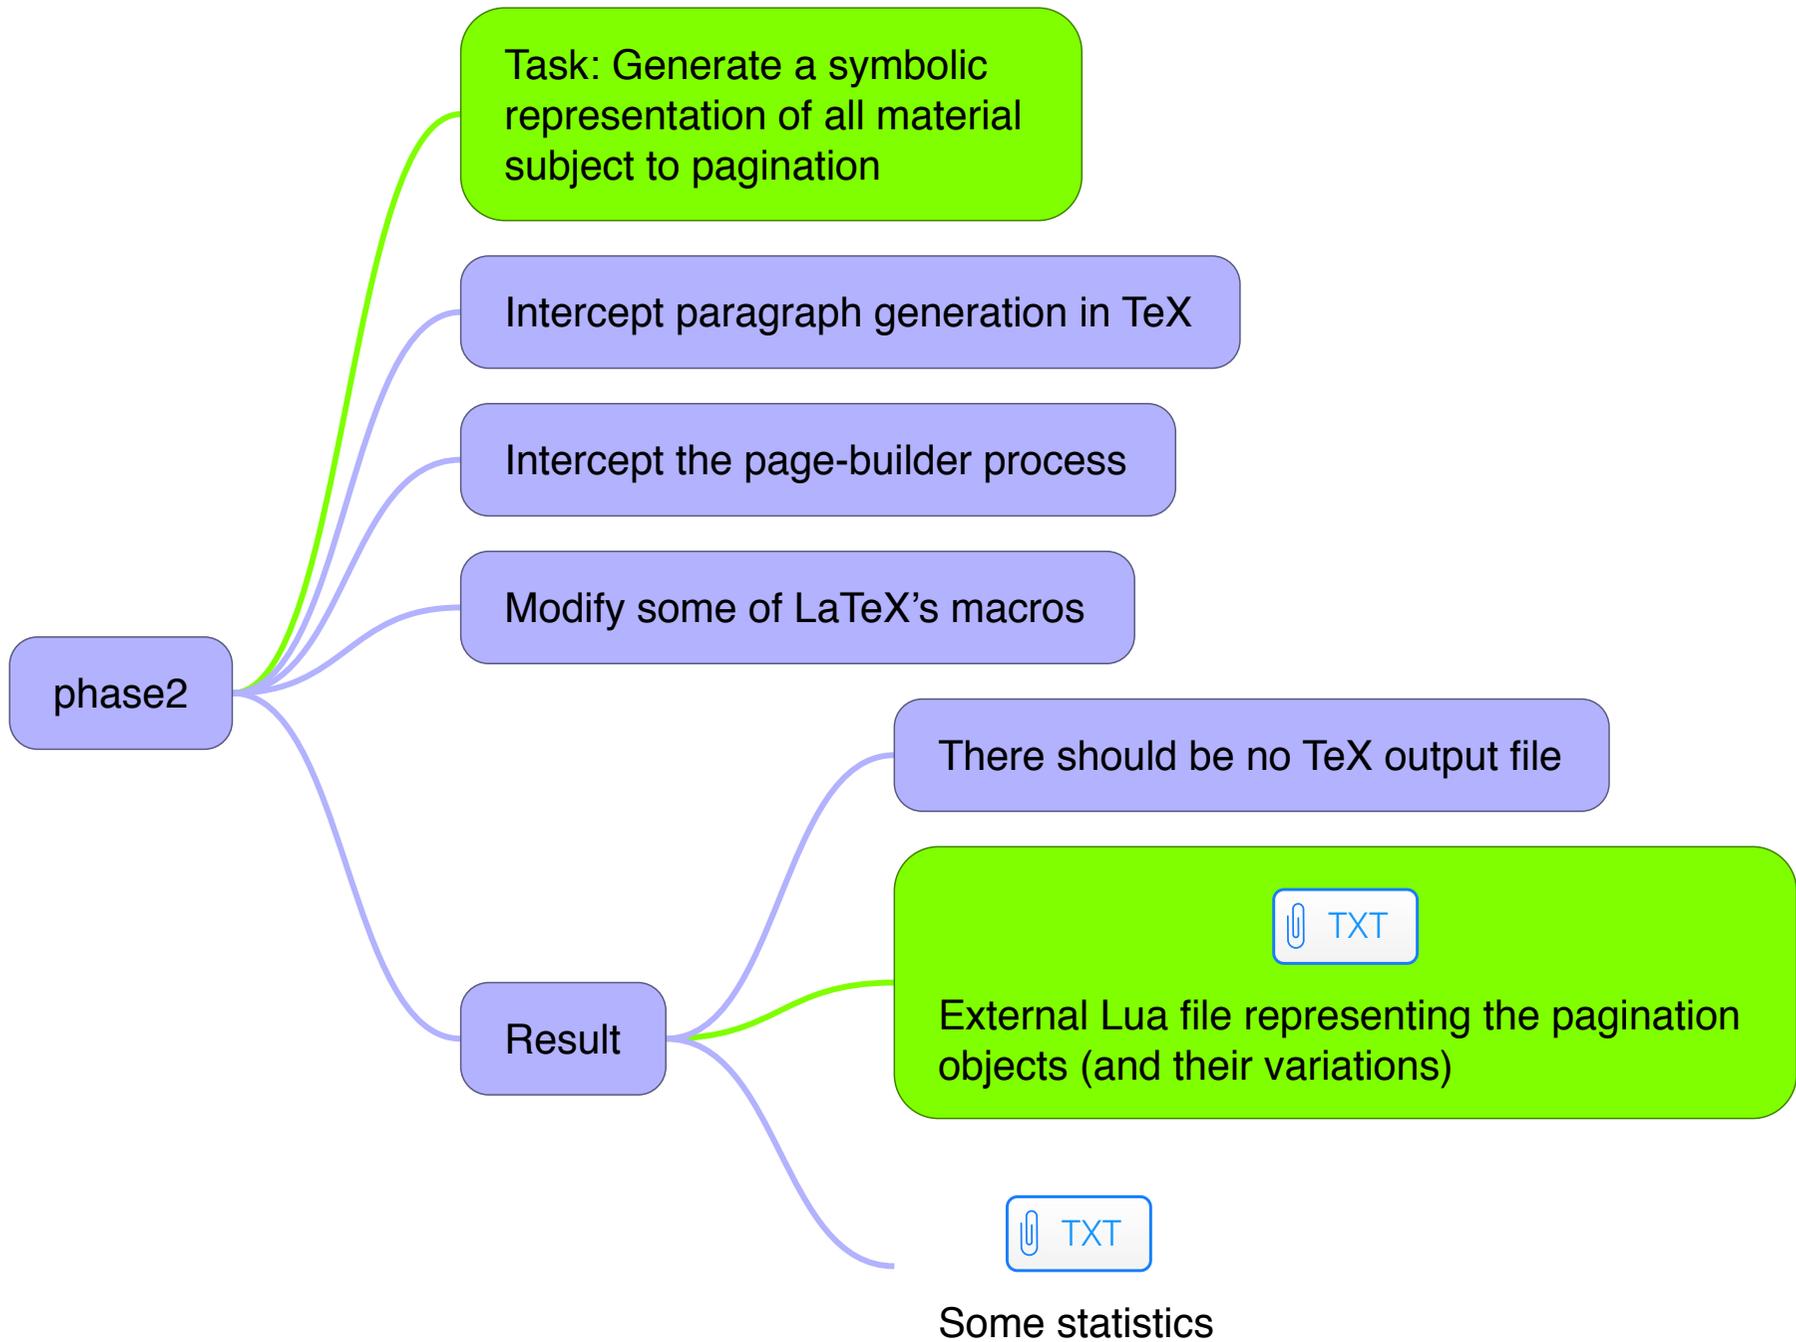
Result



Change `\clearpage` to not generate `\penalty -10001`



Change `\@starttoc` to always use TOC from phase1

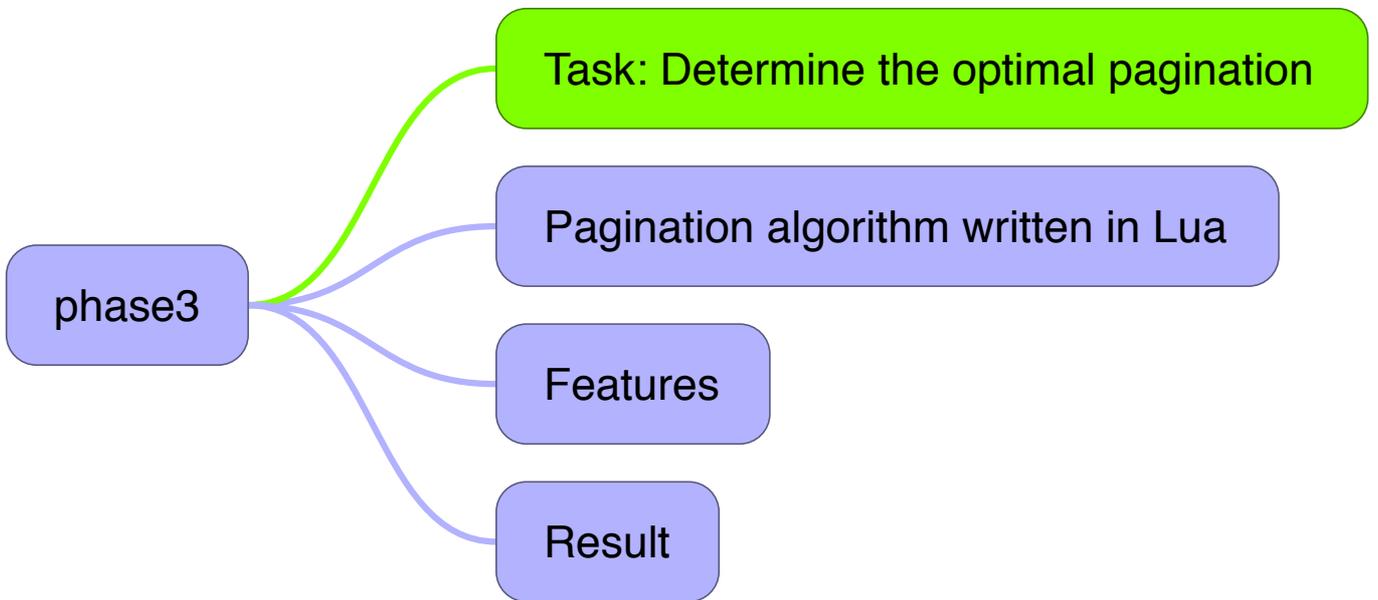


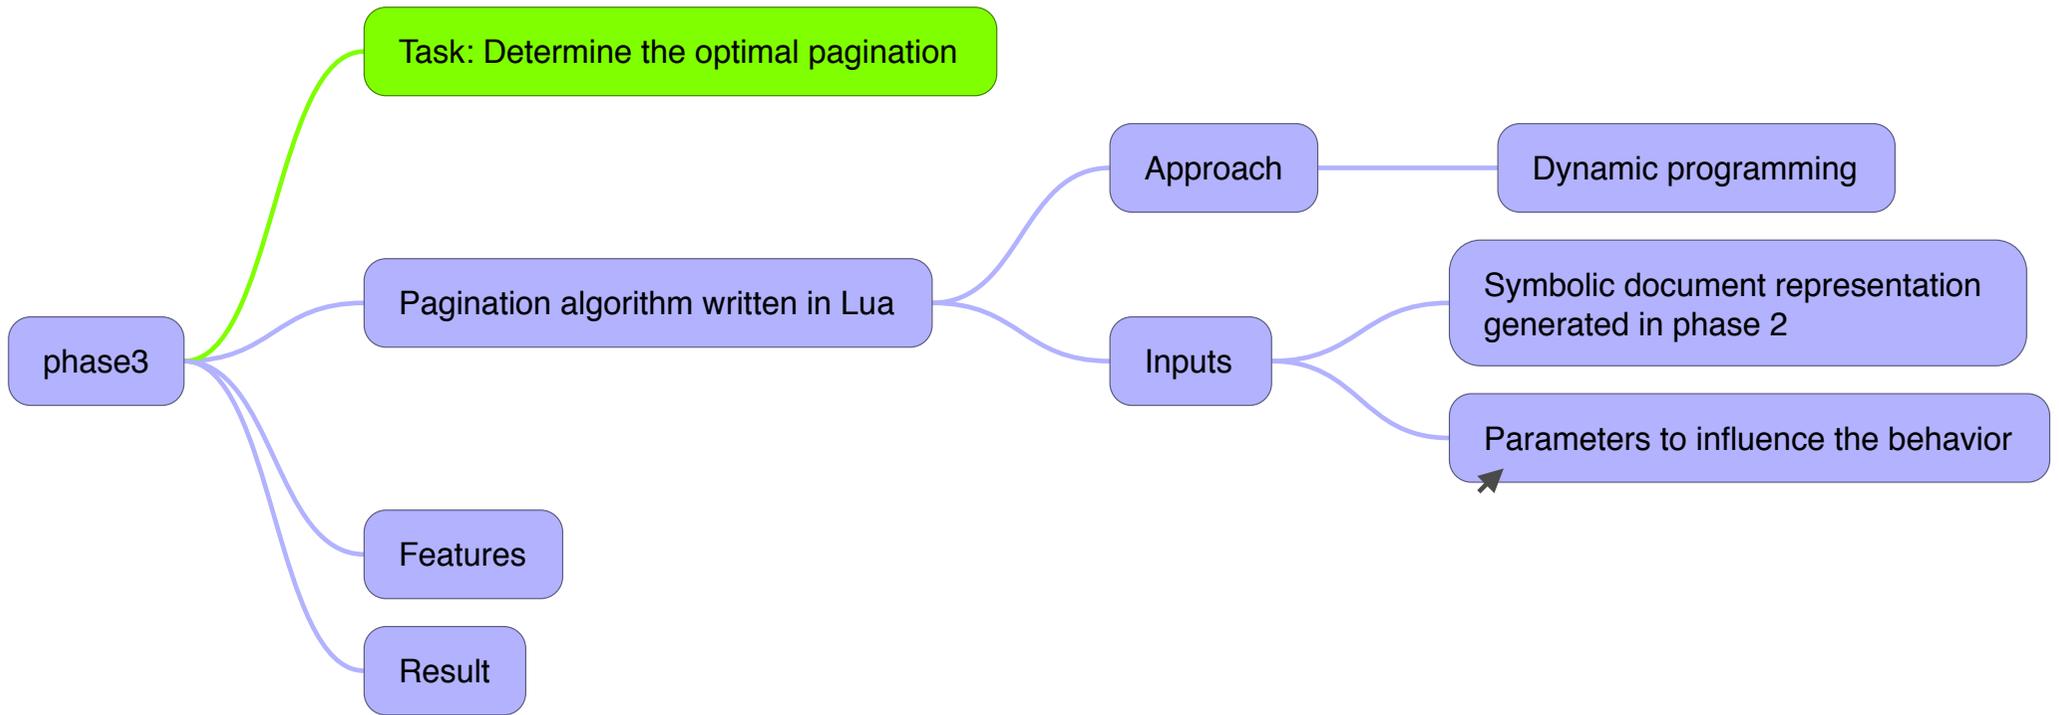


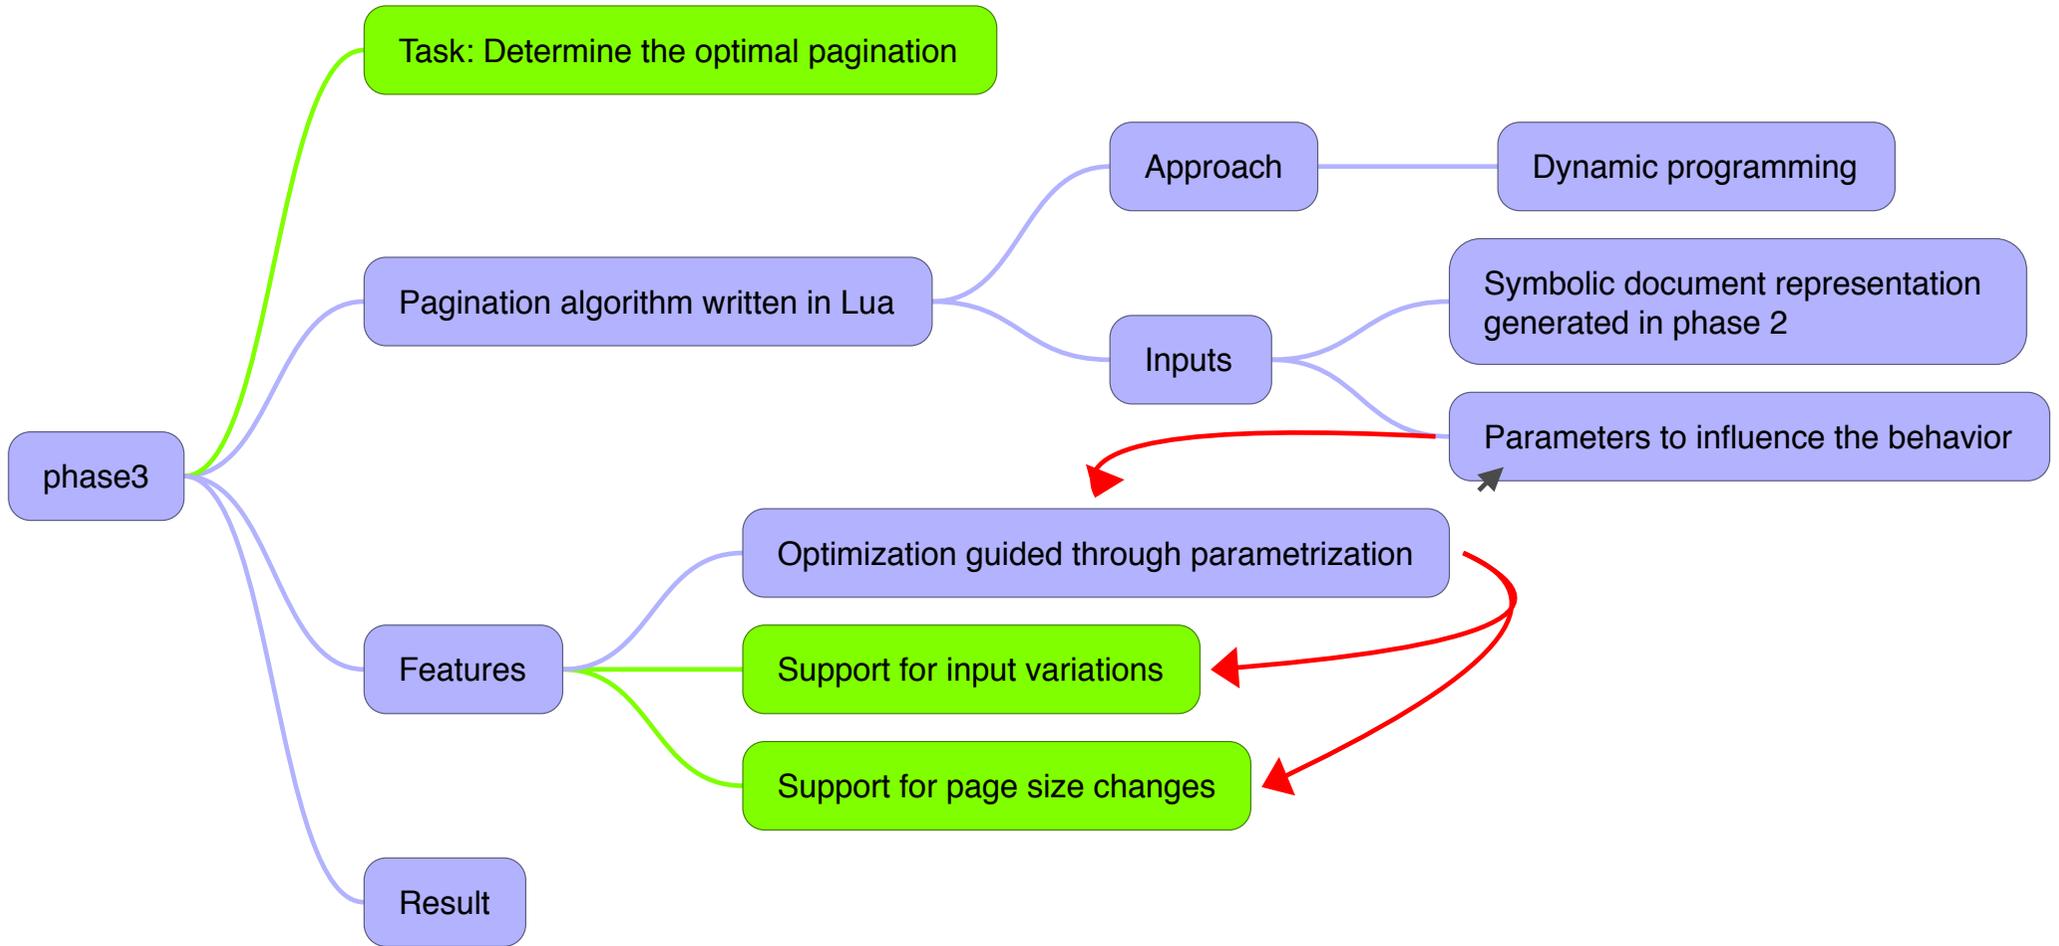
External Lua file representing the pagination objects (and their variations)

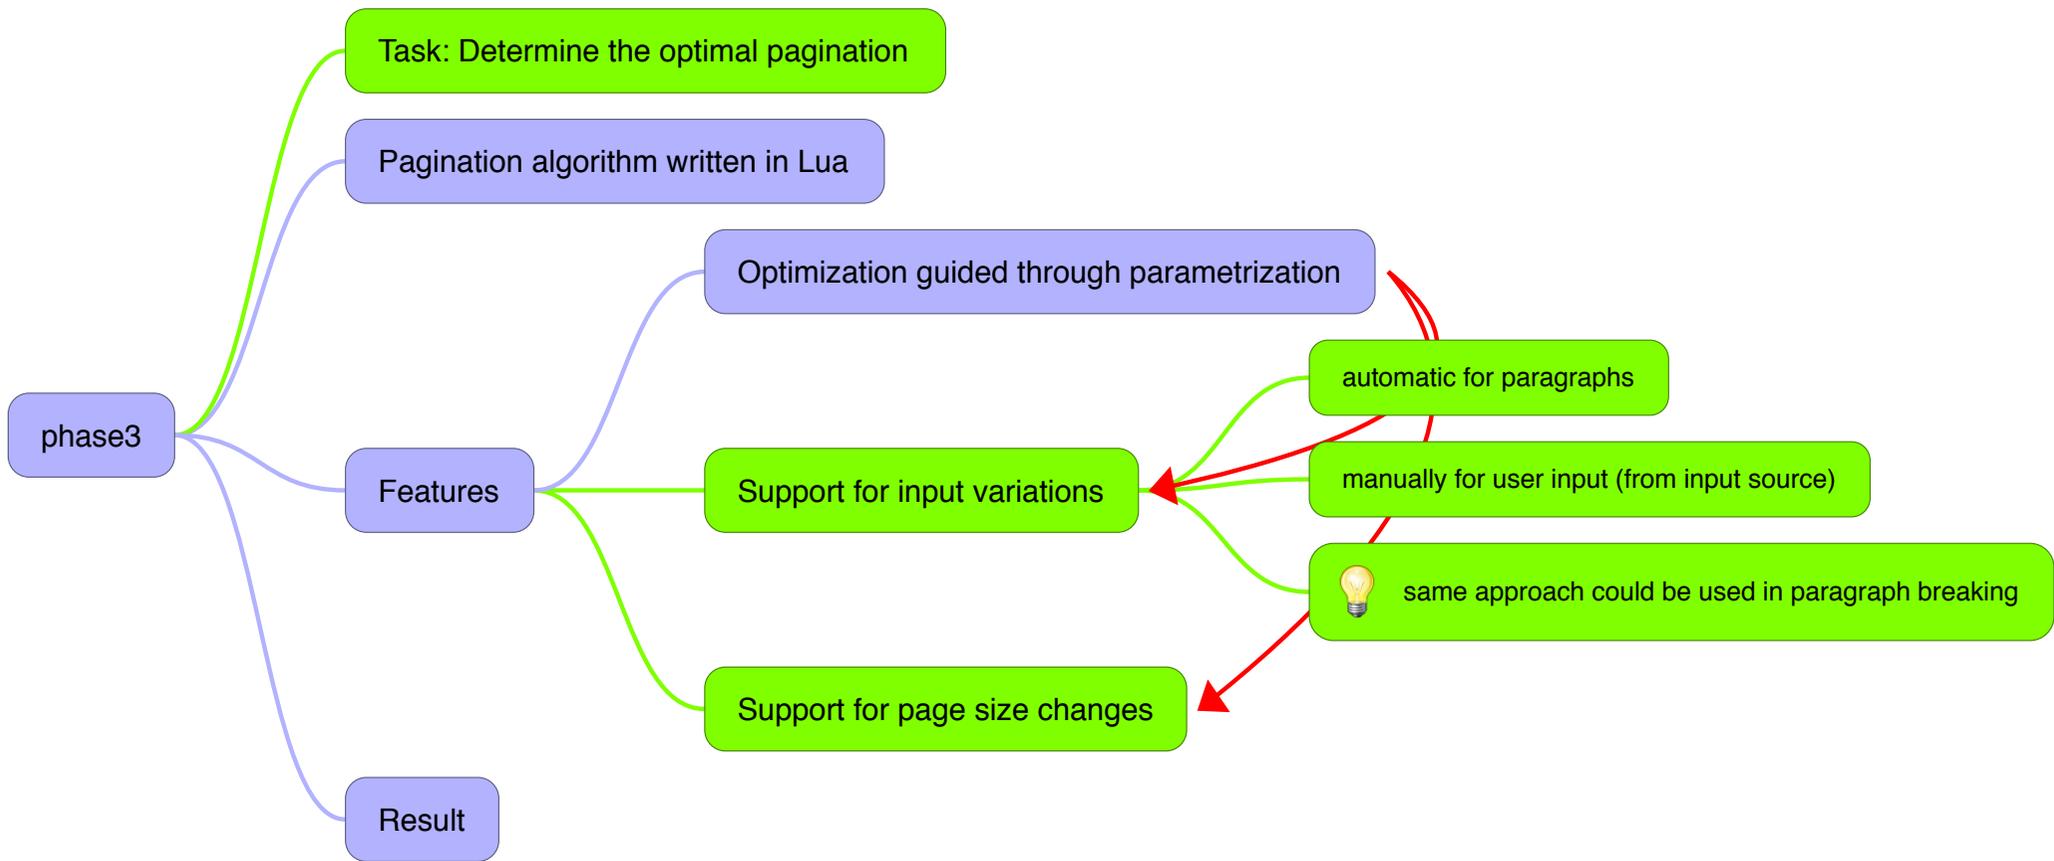


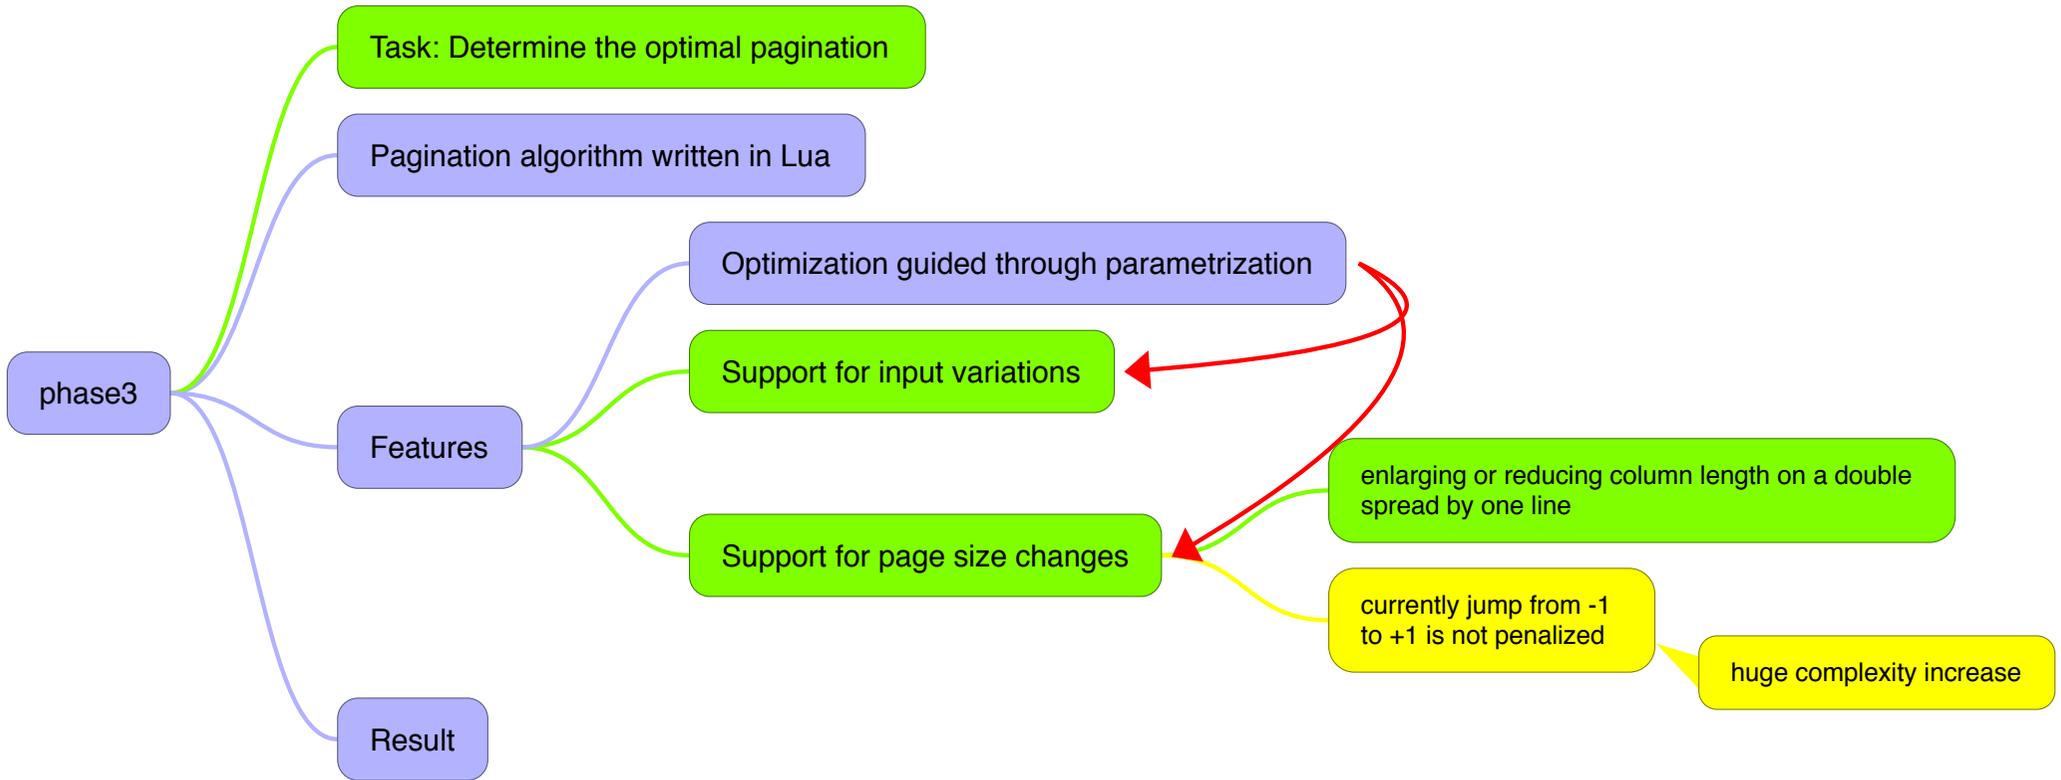
**Some statistics**

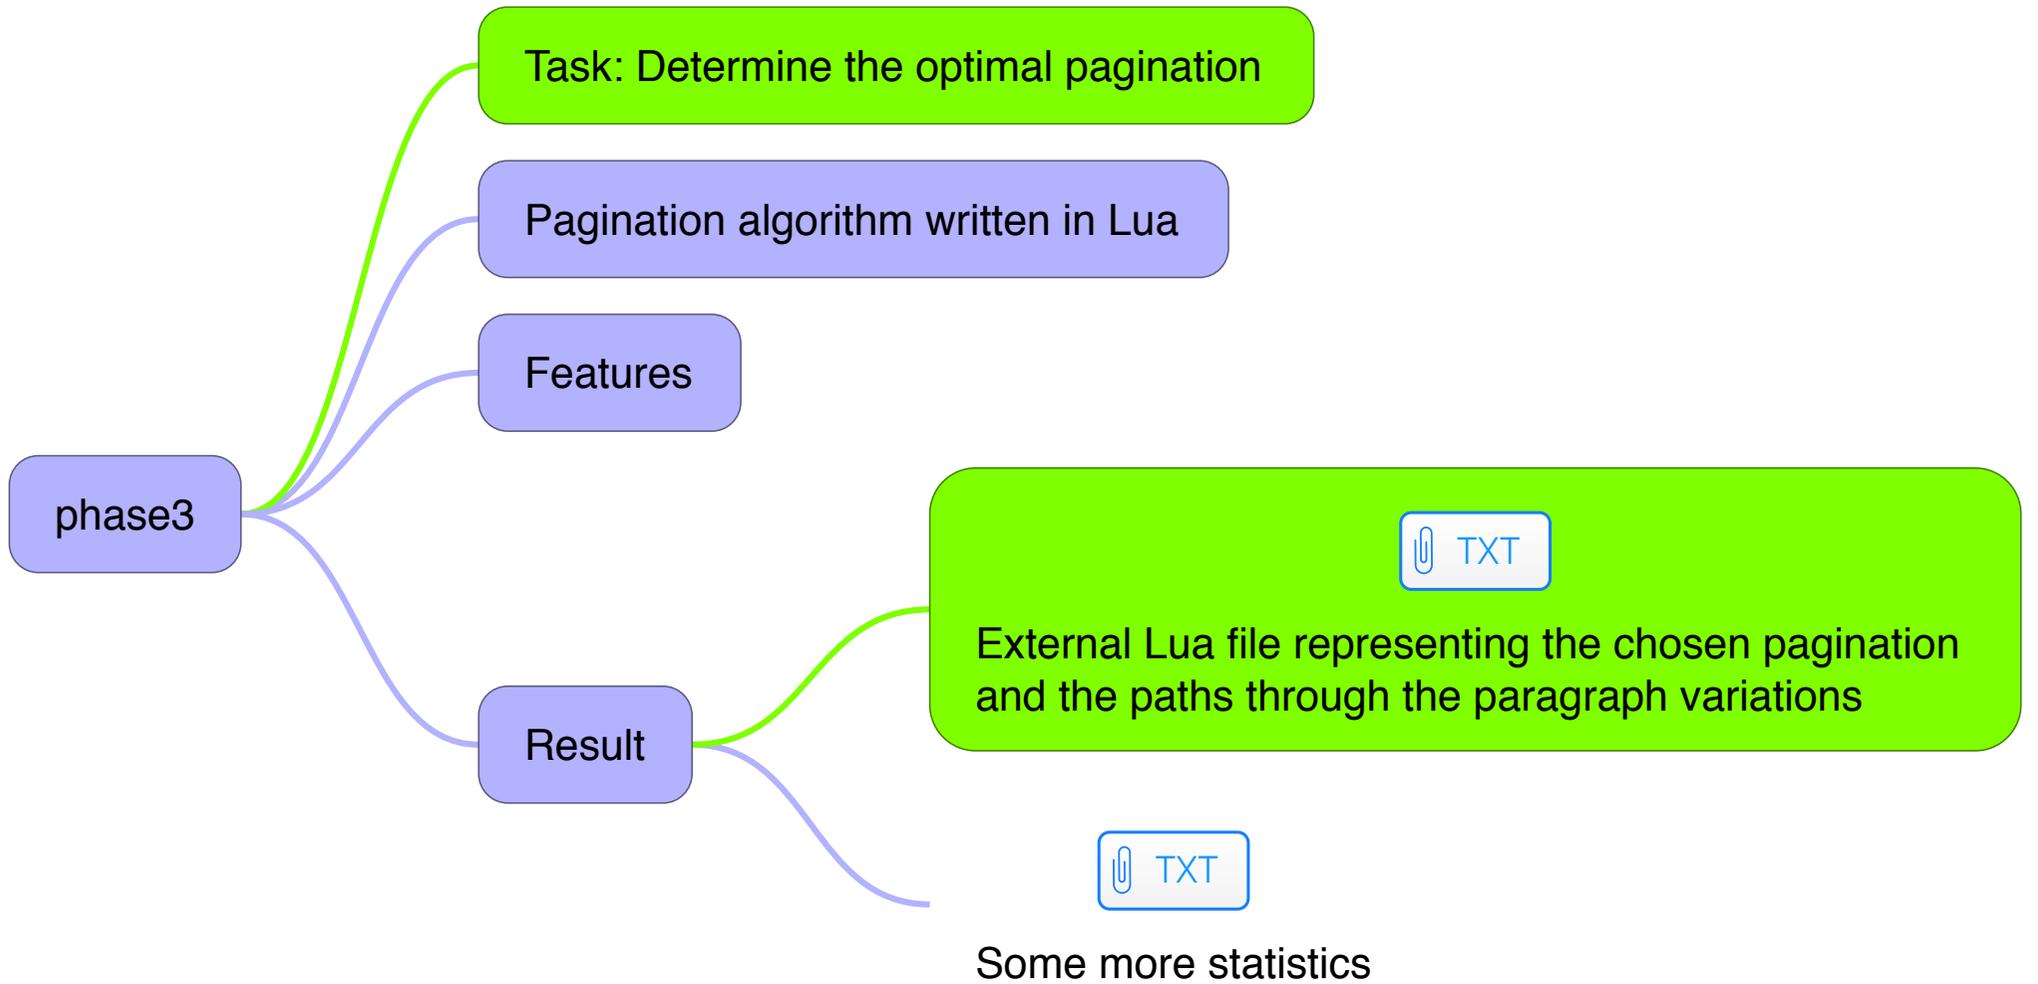












Task: Determine the optimal pagination

Pagination algorithm written in Lua

Features

phase3

Result

External Lua file representing the chosen pagination and the paths through the paragraph variations

TXT

TXT

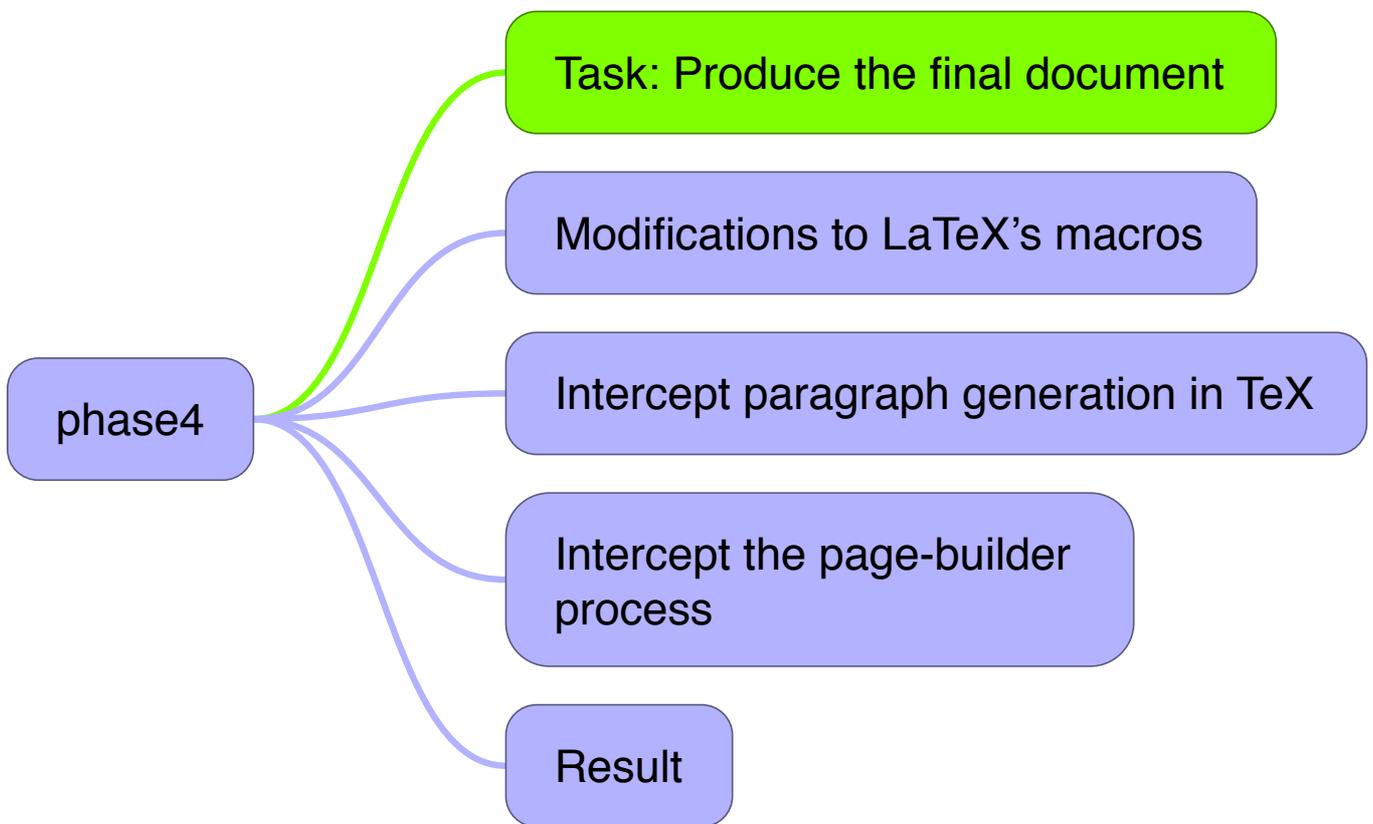
Some more statistics

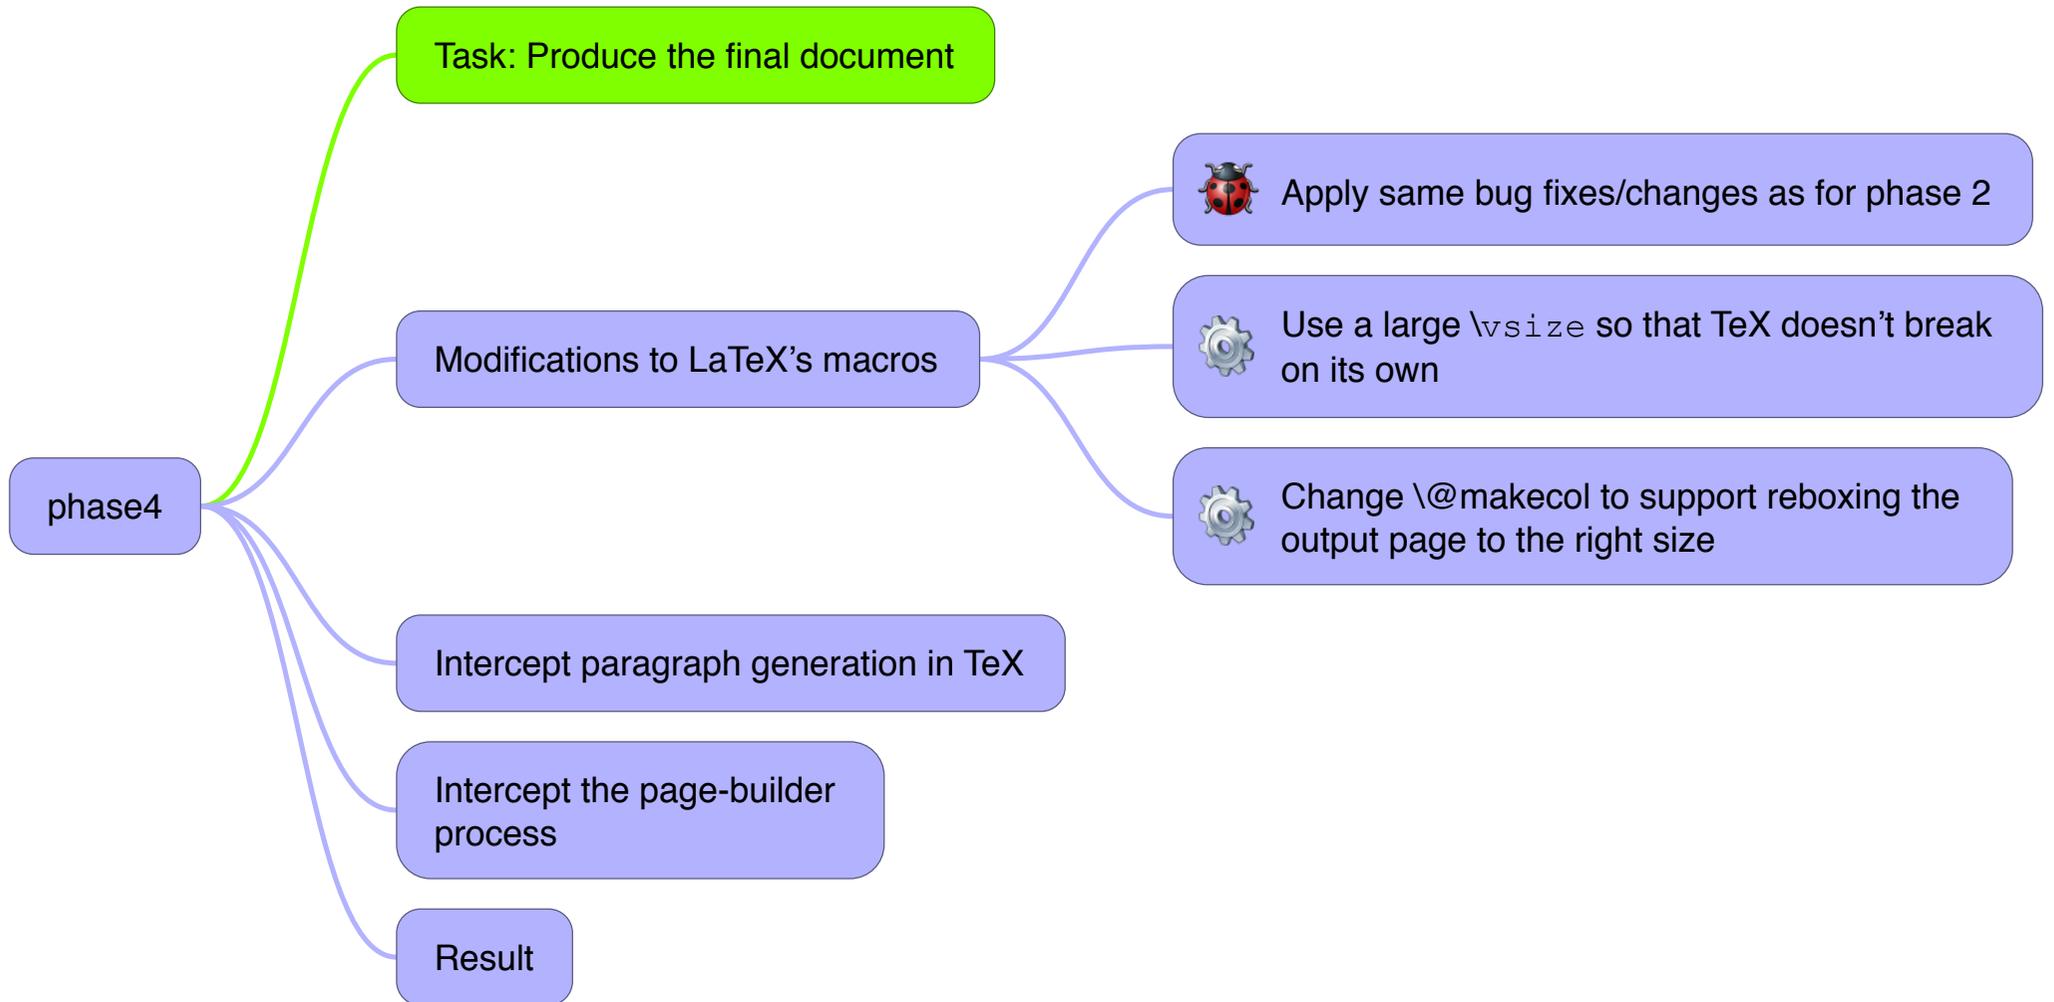


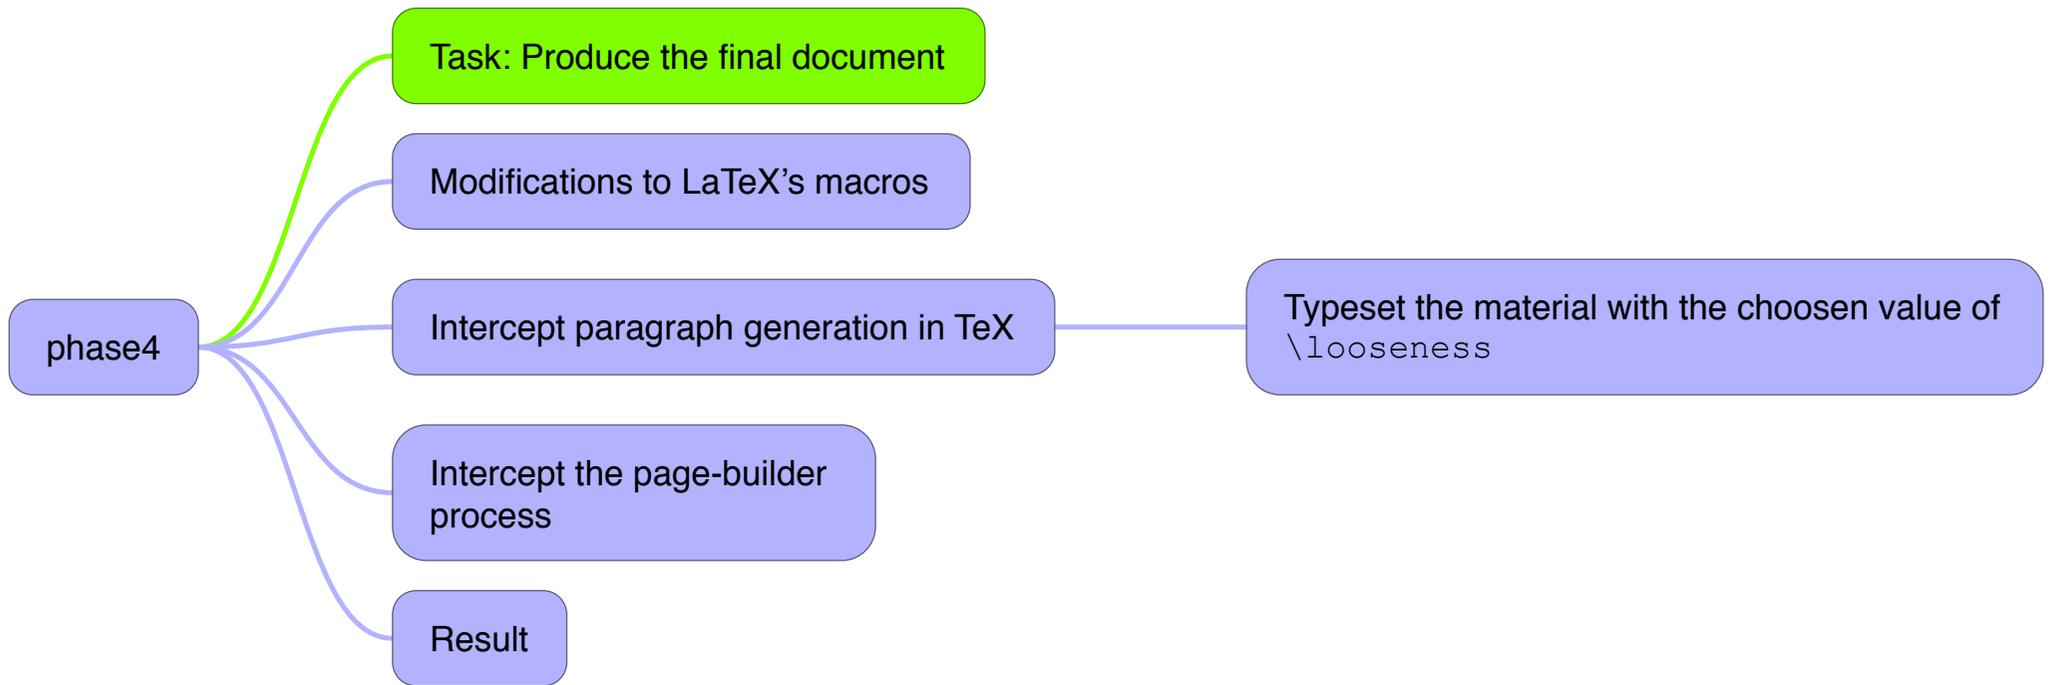
External Lua file representing the chosen pagination and the paths through the paragraph variations

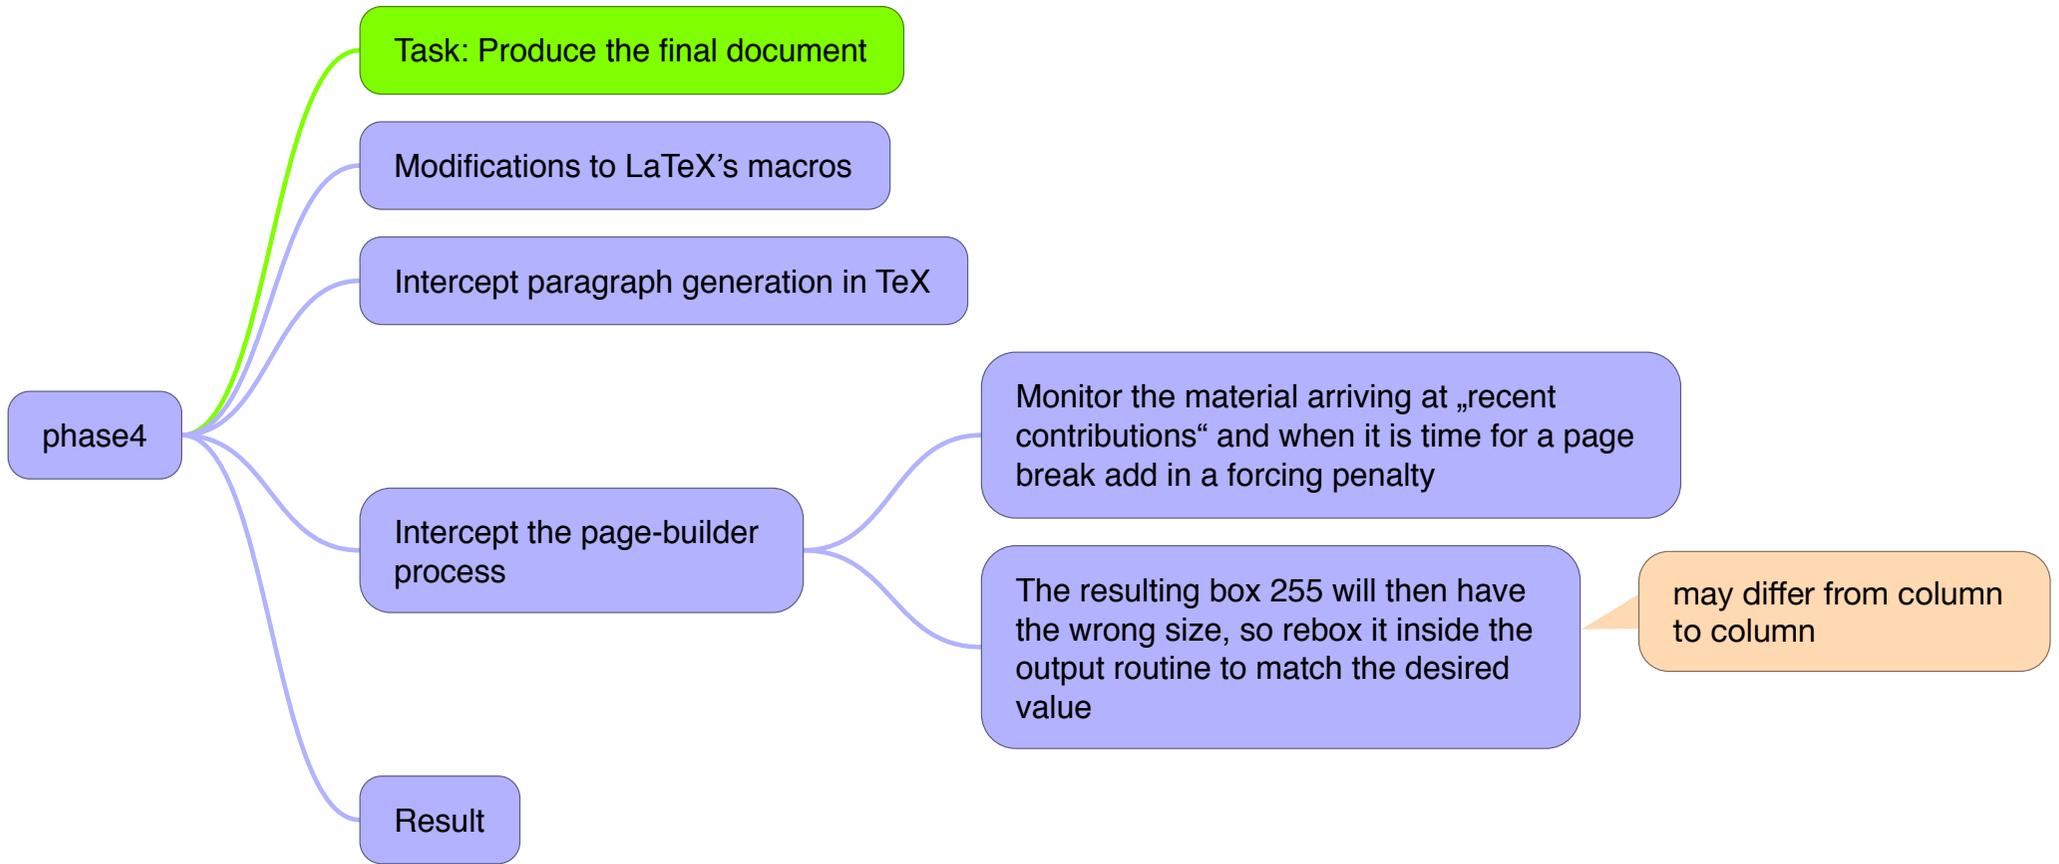


**Some more statistics**









phase4

Task: Produce the final document

Modifications to LaTeX's macros

Intercept paragraph generation in TeX

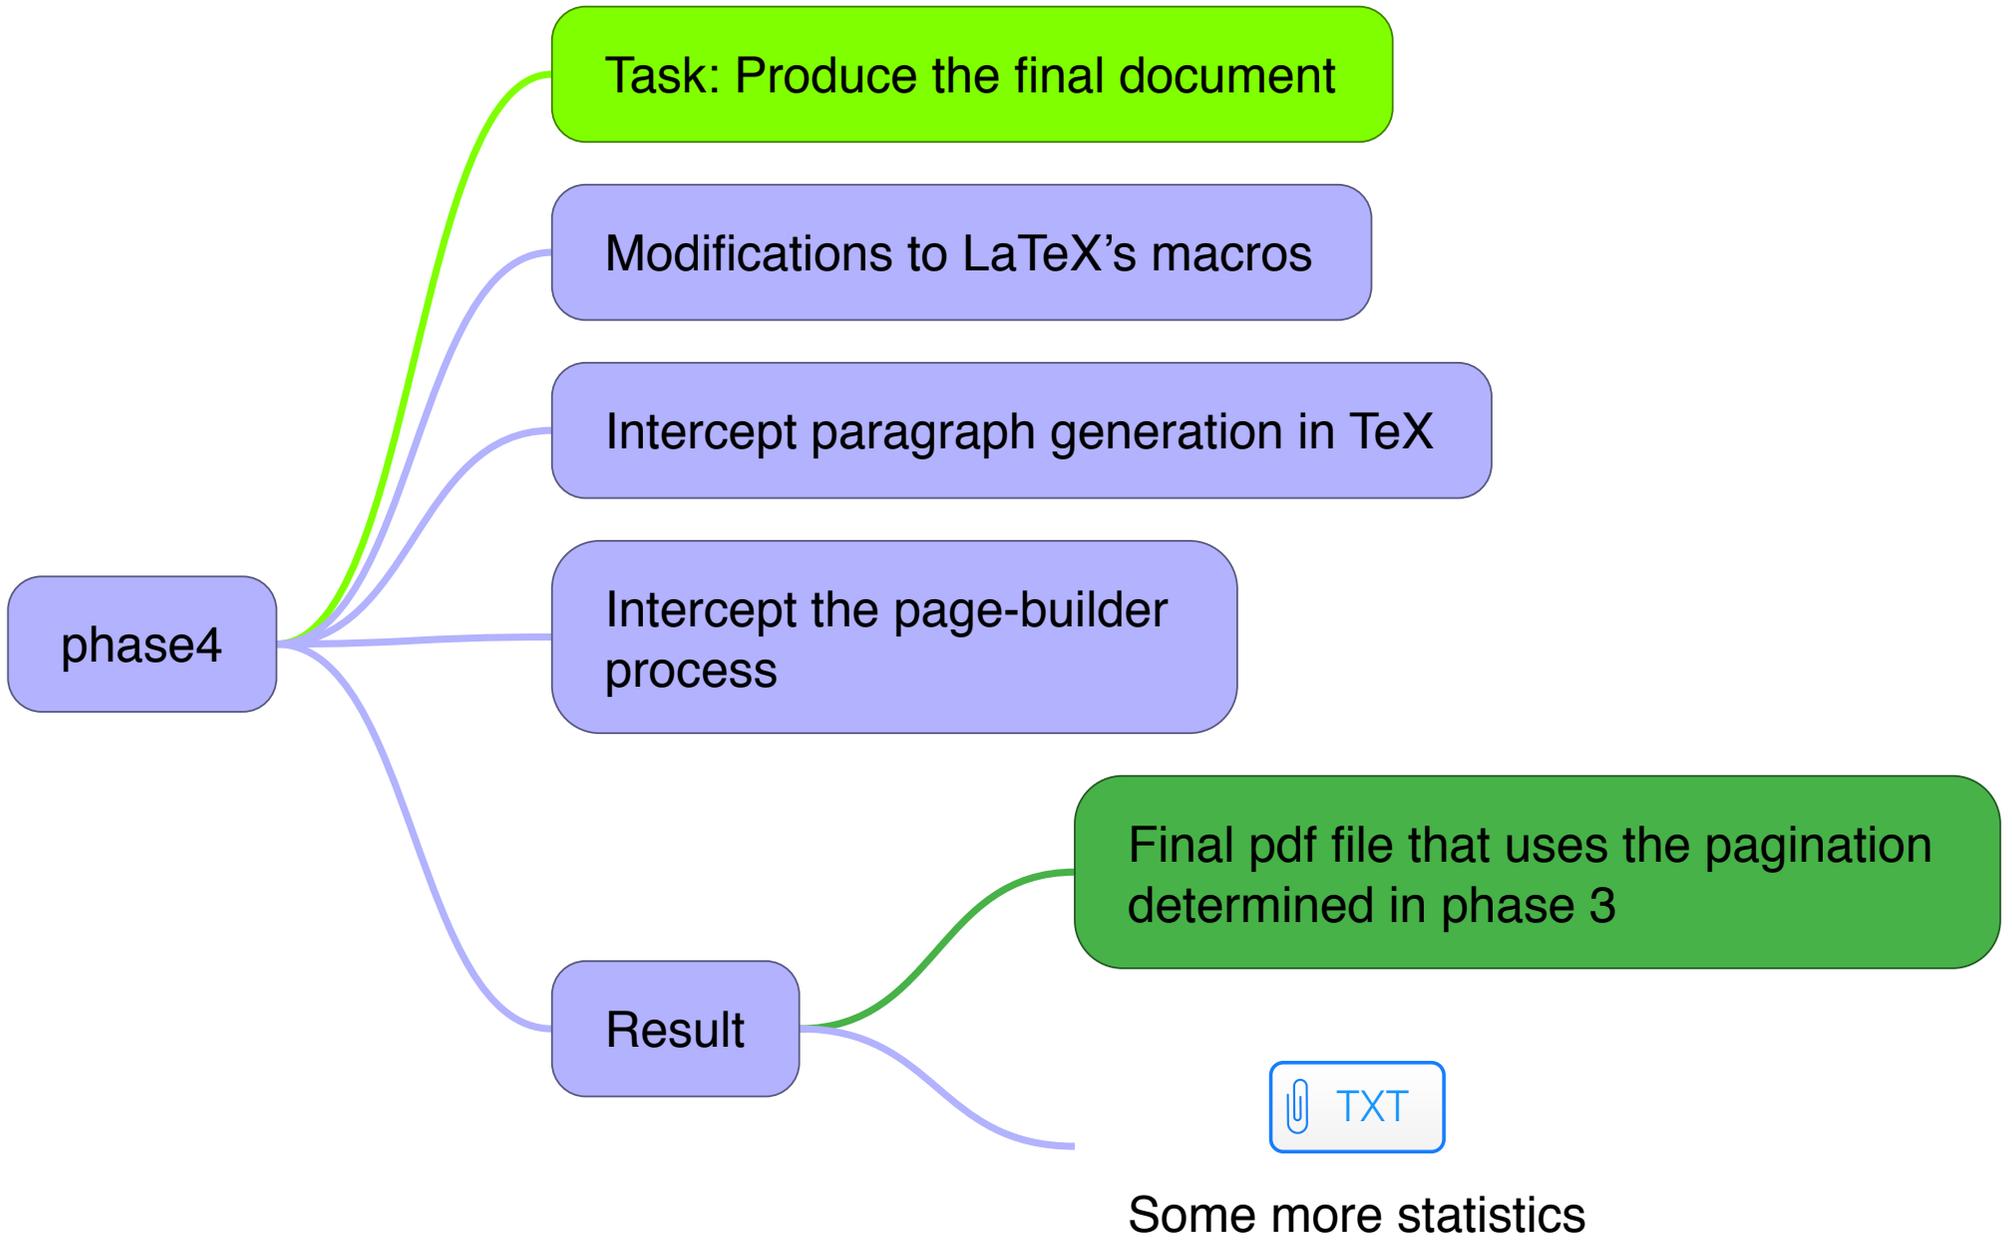
Intercept the page-builder process

Result

Monitor the material arriving at „recent contributions“ and when it is time for a page break add in a forcing penalty

The resulting box 255 will then have the wrong size, so rebox it inside the output routine to match the desired value

may differ from column to column



Task: Produce the final document

Modifications to LaTeX's macros

Intercept paragraph generation in TeX

Intercept the page-builder process

phase4

Result

Final pdf file that uses the pagination determined in phase 3



Some more statistics



**Some more statistics**



**First results**



Typesetting Alice again

Typesetting Grimm Fairy Tales

Typesetting Pride and Prejudice  
by Jane Austen



Typesetting Alice again

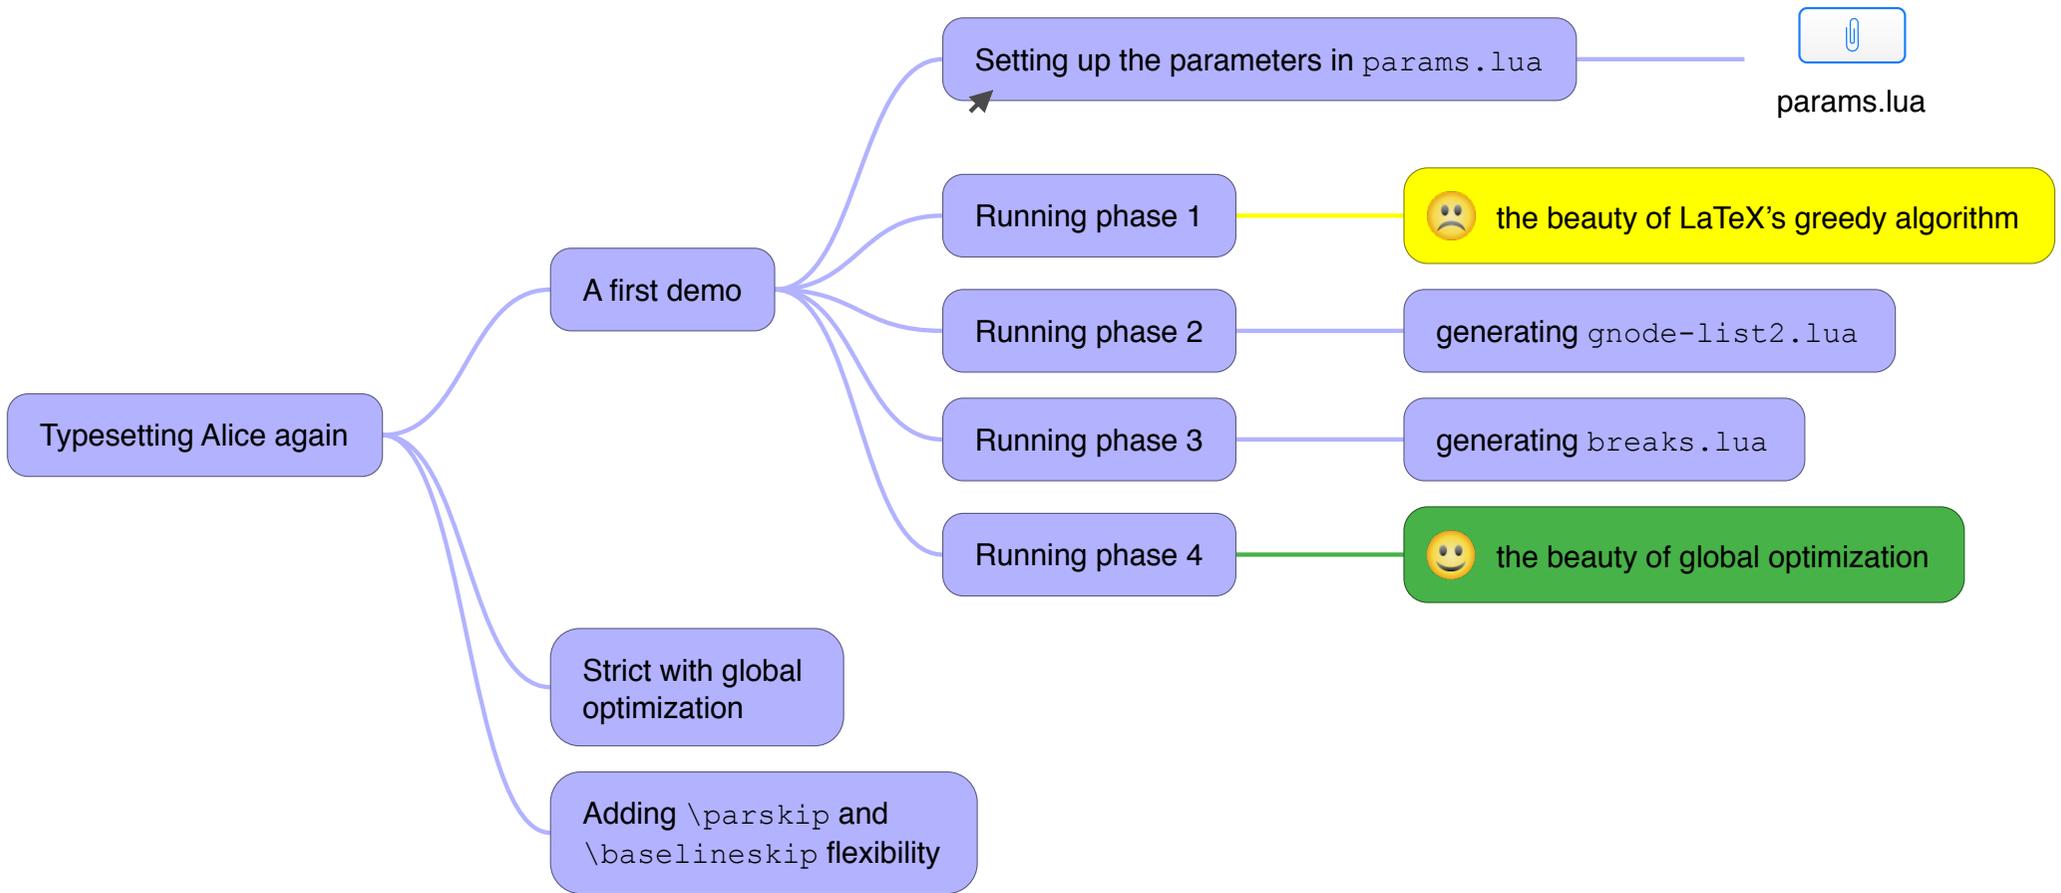
A first demo

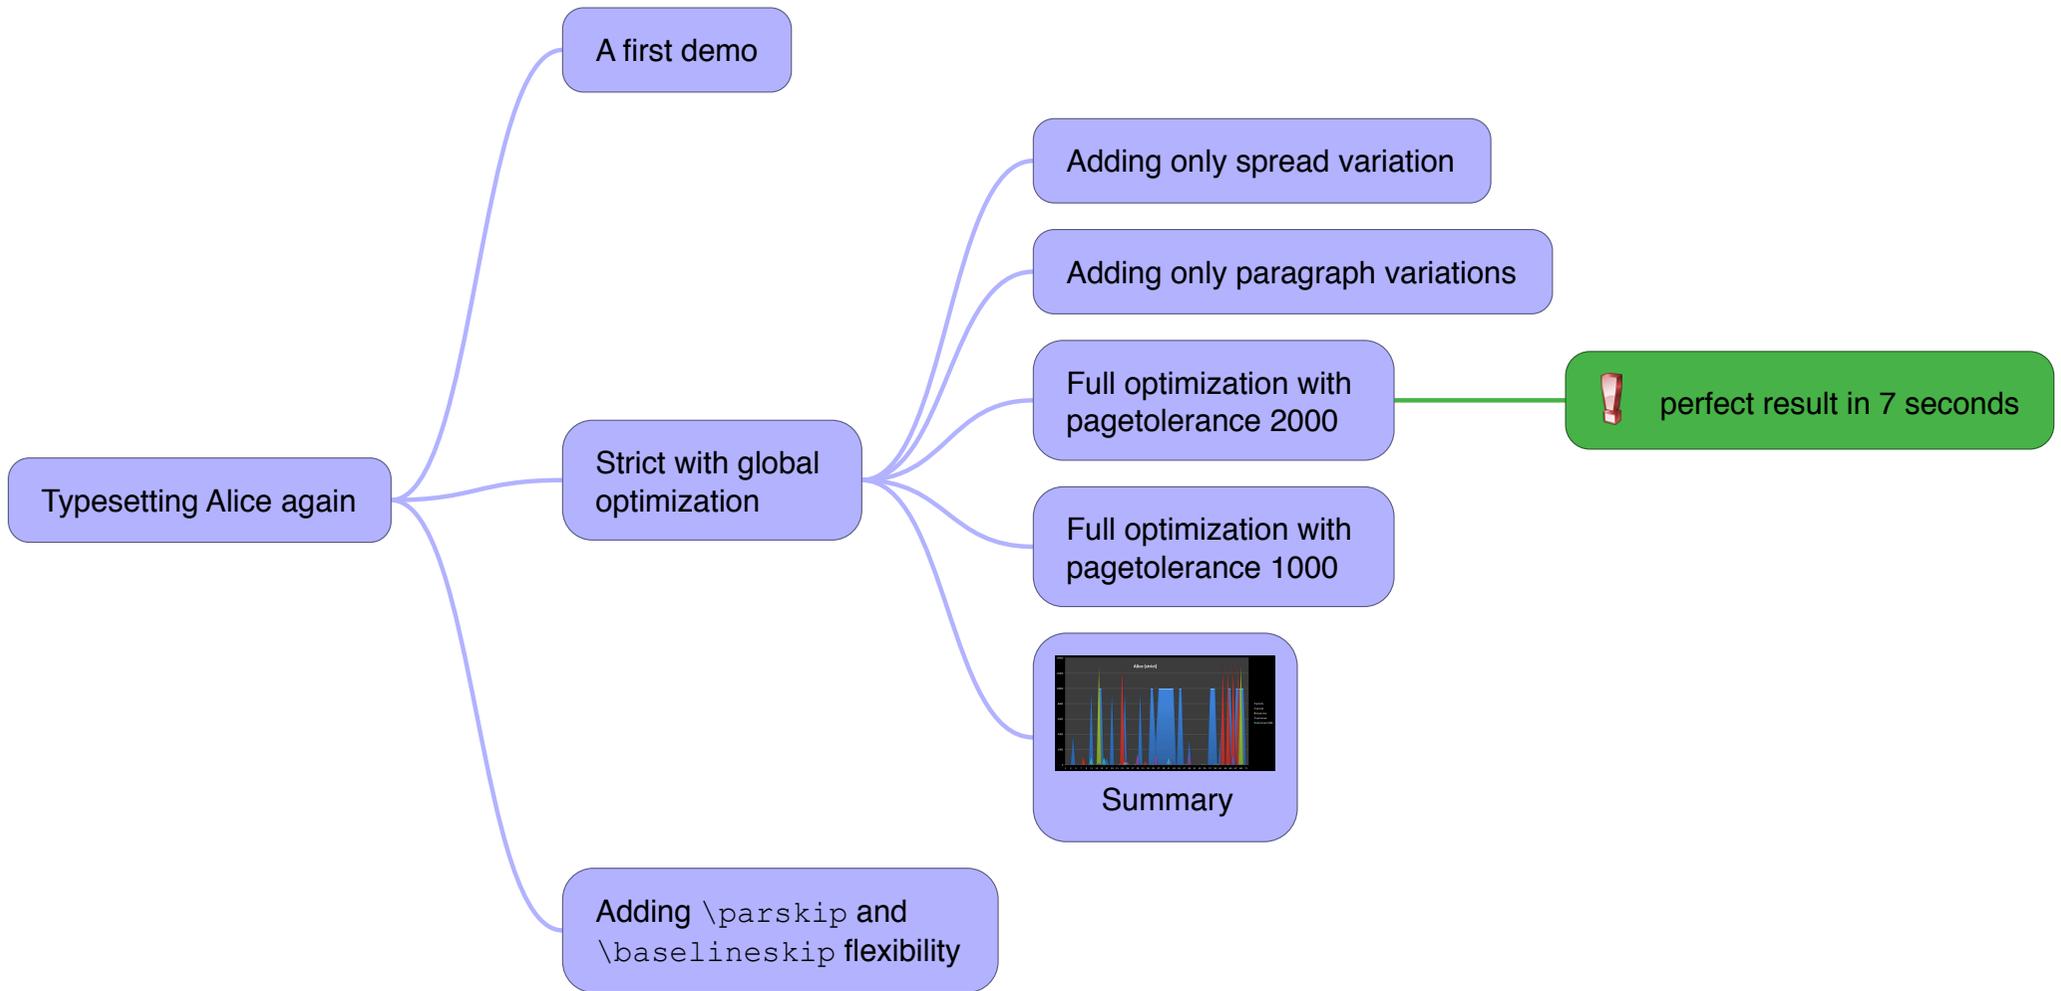
Strict with global optimization

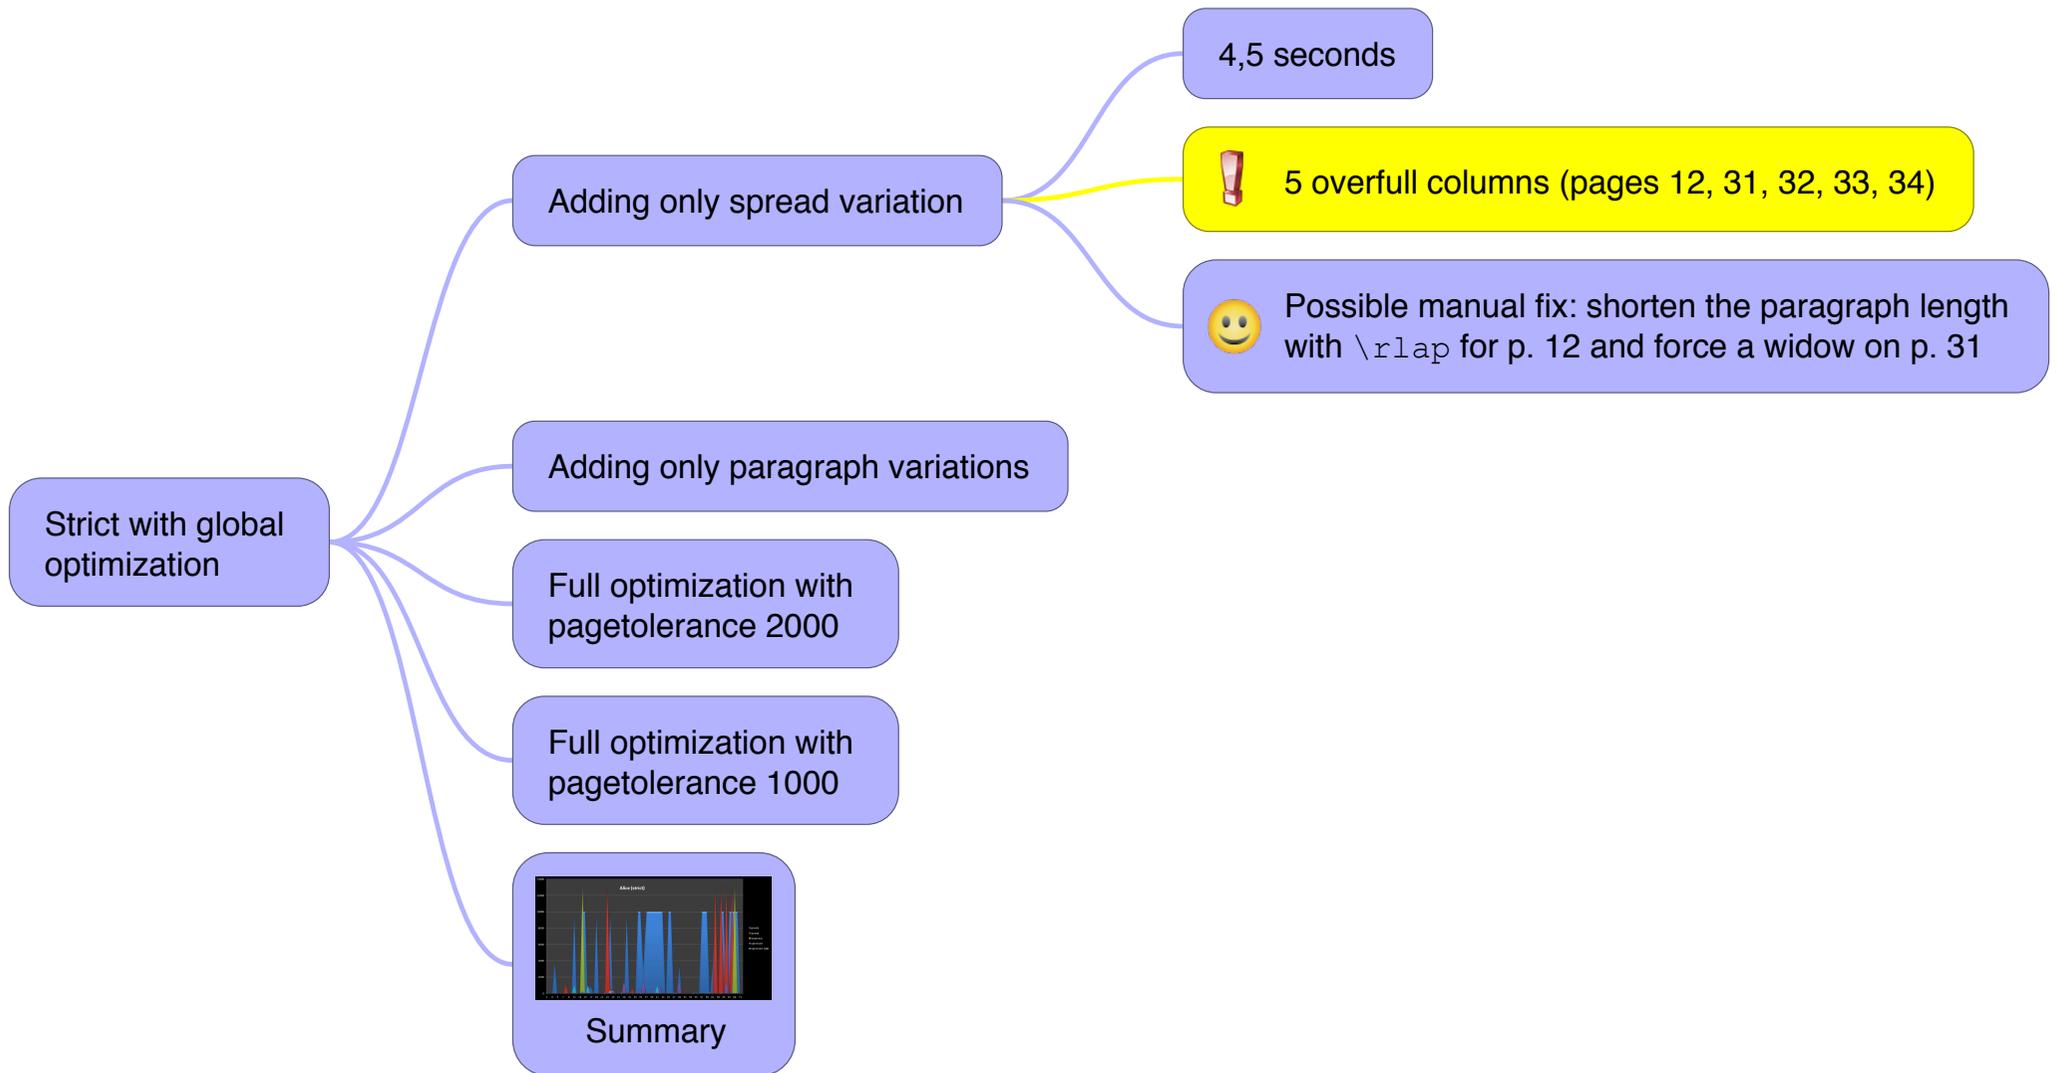
Adding `\parskip` and `\baselineskip` flexibility

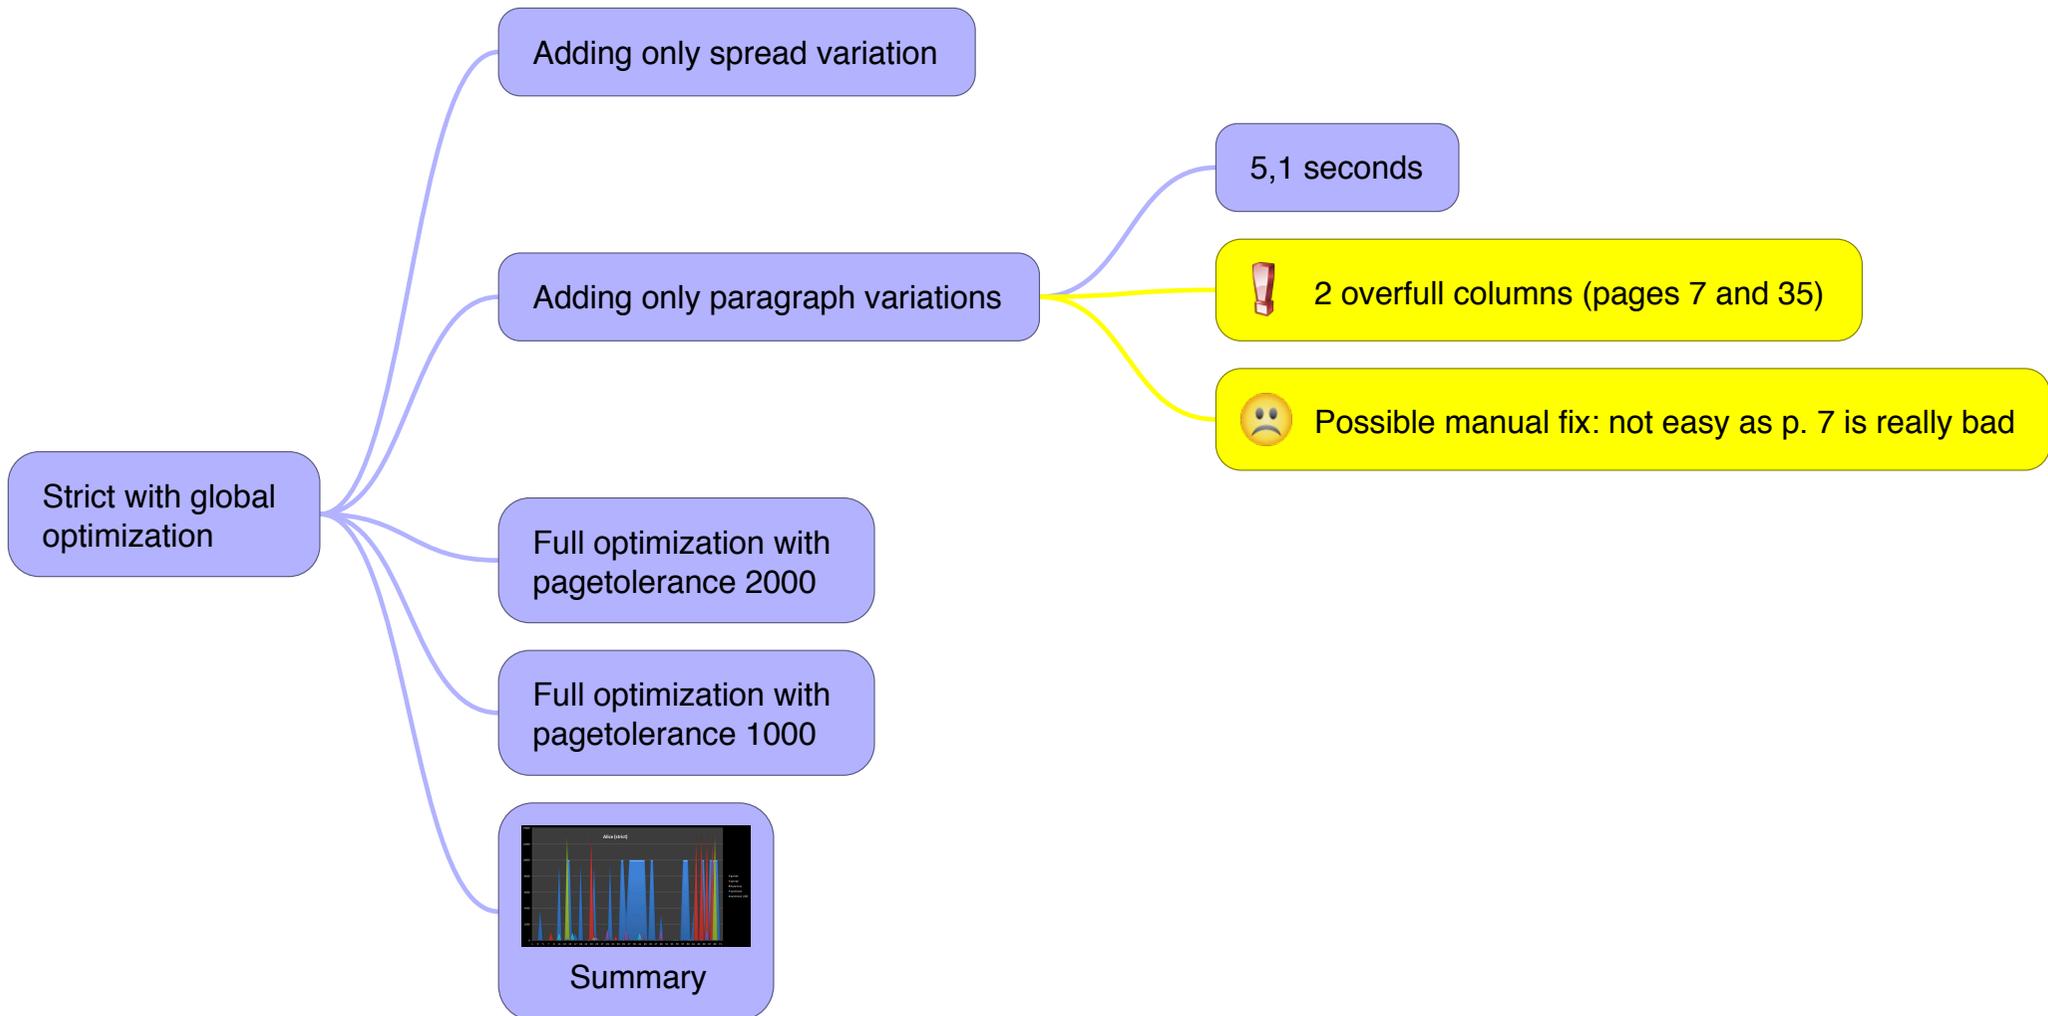
Typesetting Grimm Fairy Tales

Typesetting Pride and Prejudice by Jane Austen









Strict with global optimization

Adding only spread variation

Adding only paragraph variations

Full optimization with pagetolerance 2000

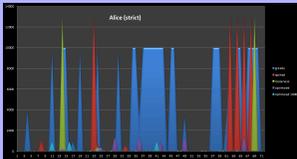
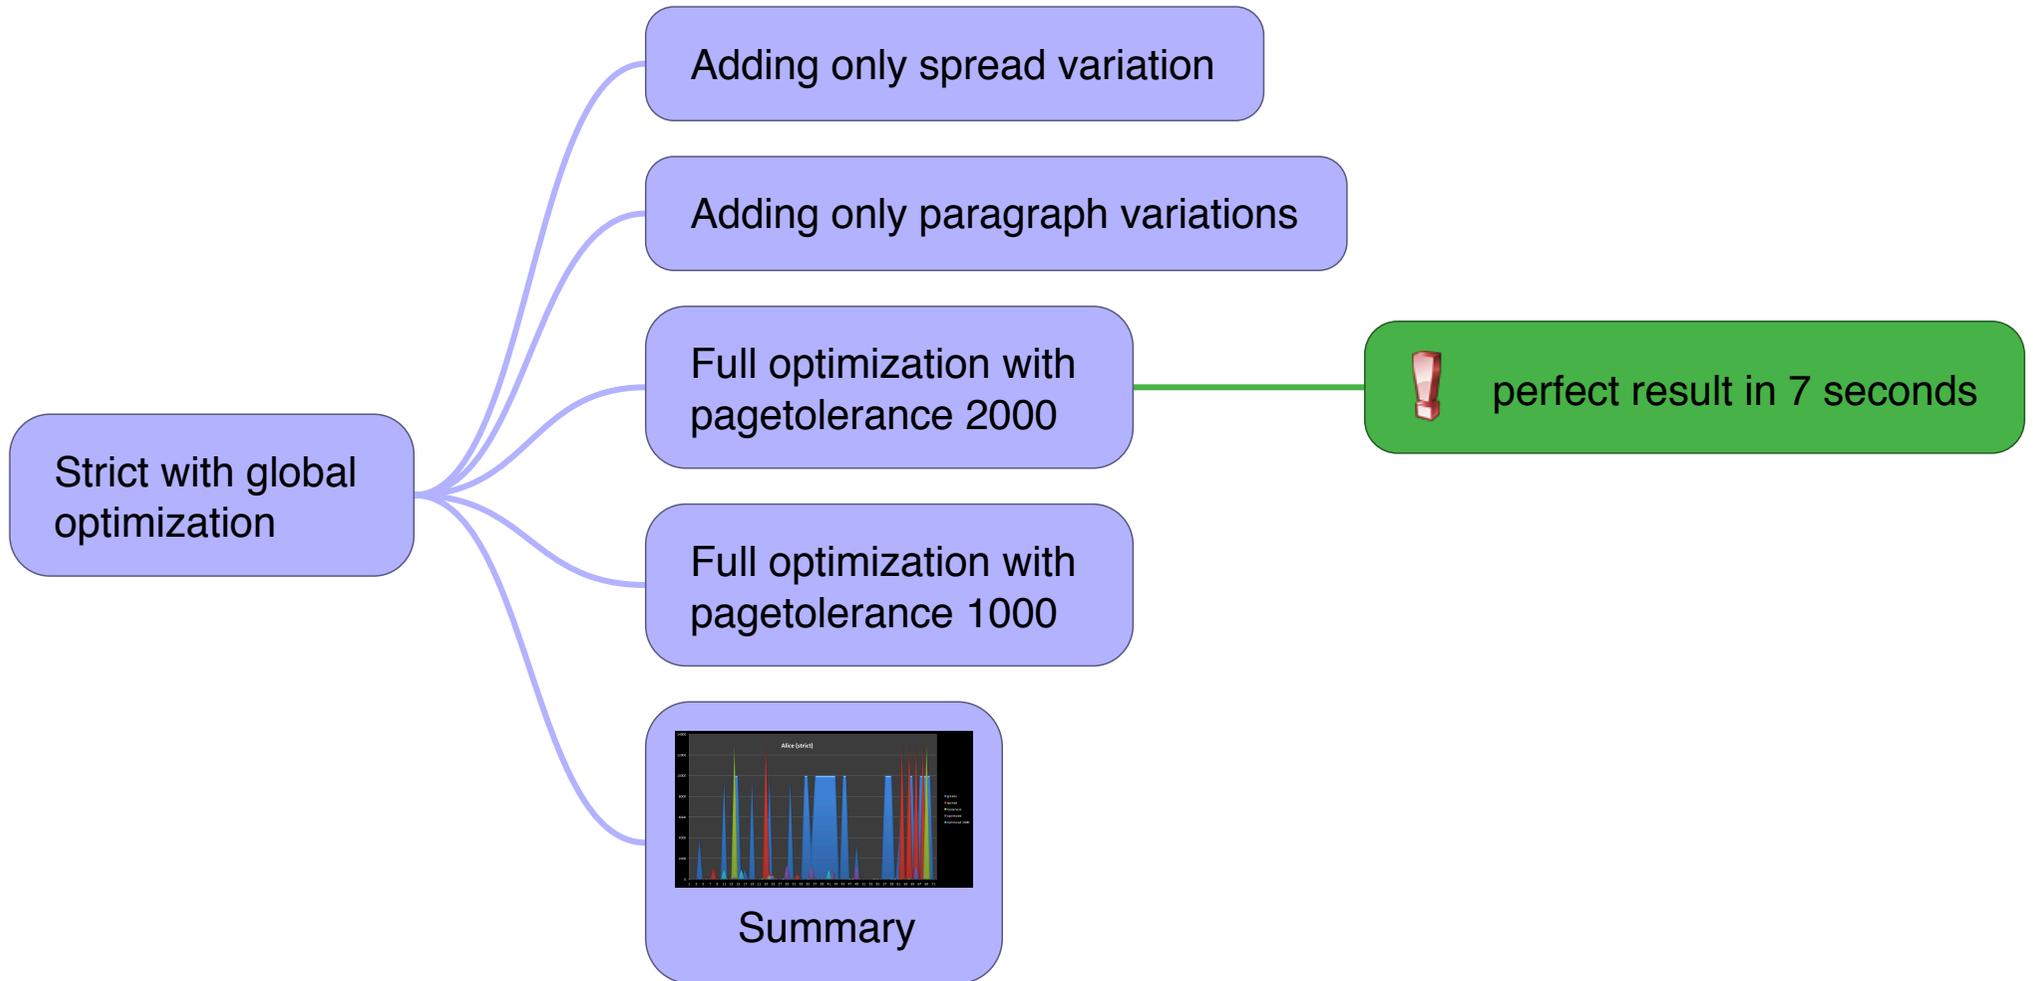
Full optimization with pagetolerance 1000

Summary

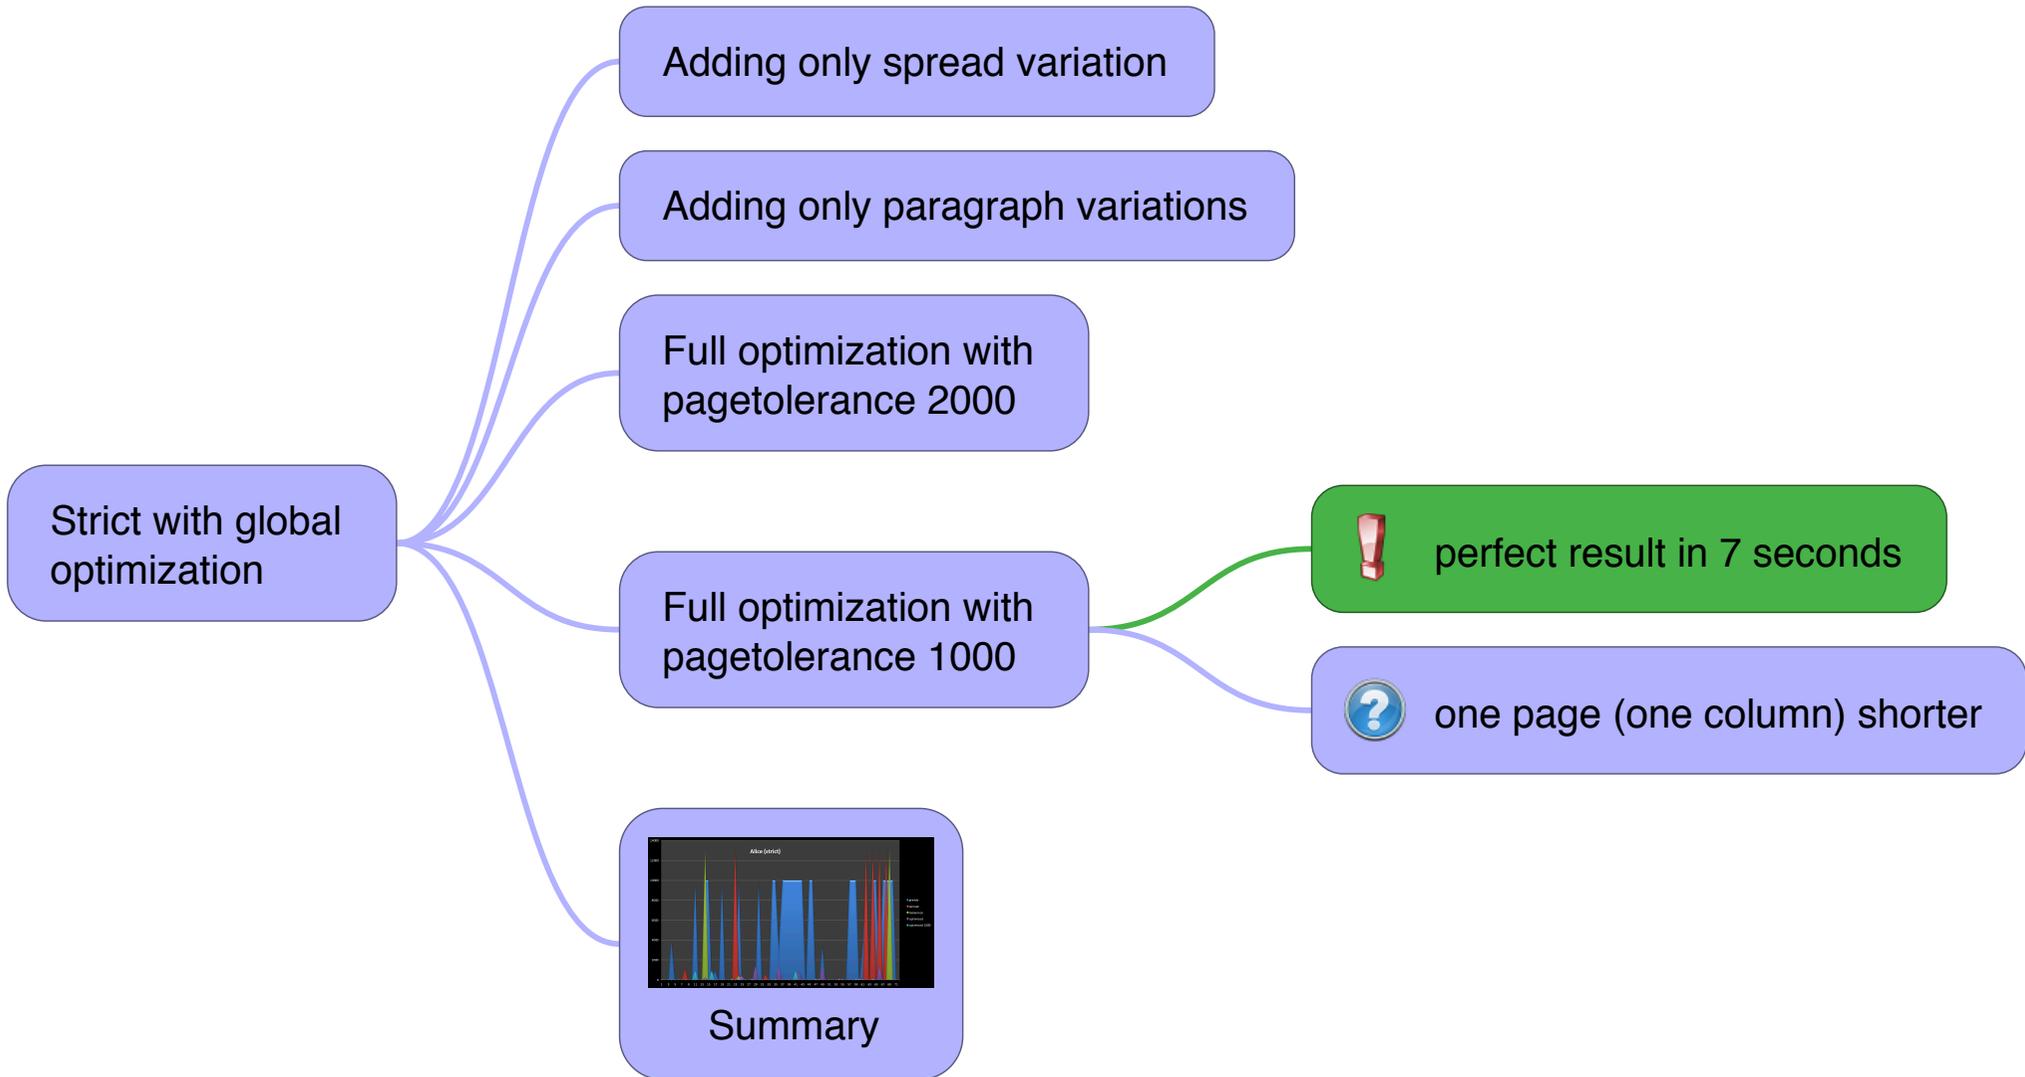
5,1 seconds

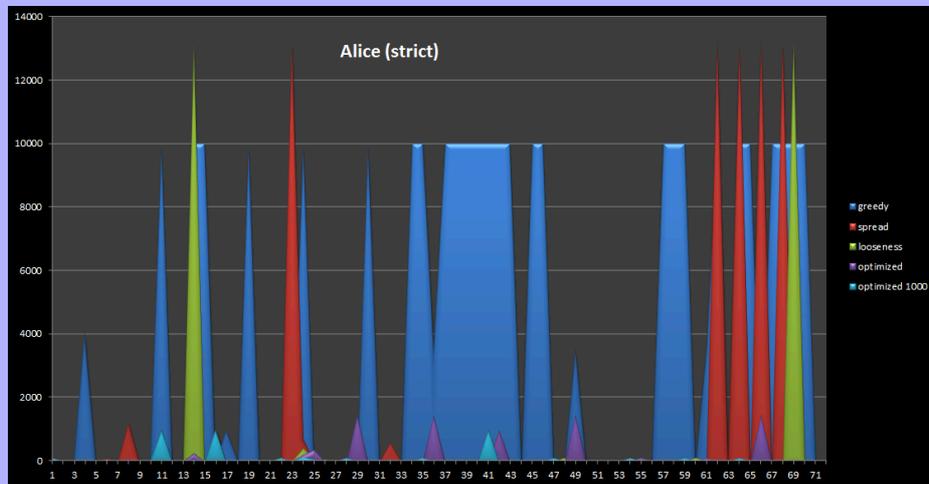
! 2 overfull columns (pages 7 and 35)

😞 Possible manual fix: not easy as p. 7 is really bad



Summary





# Summary

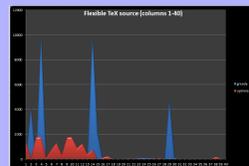
Adding `\parskip` and  
`\baselineskip` flexibility

Not really a great idea, as varying  
`\baselineskip` changes the grey  
value of the paragraph

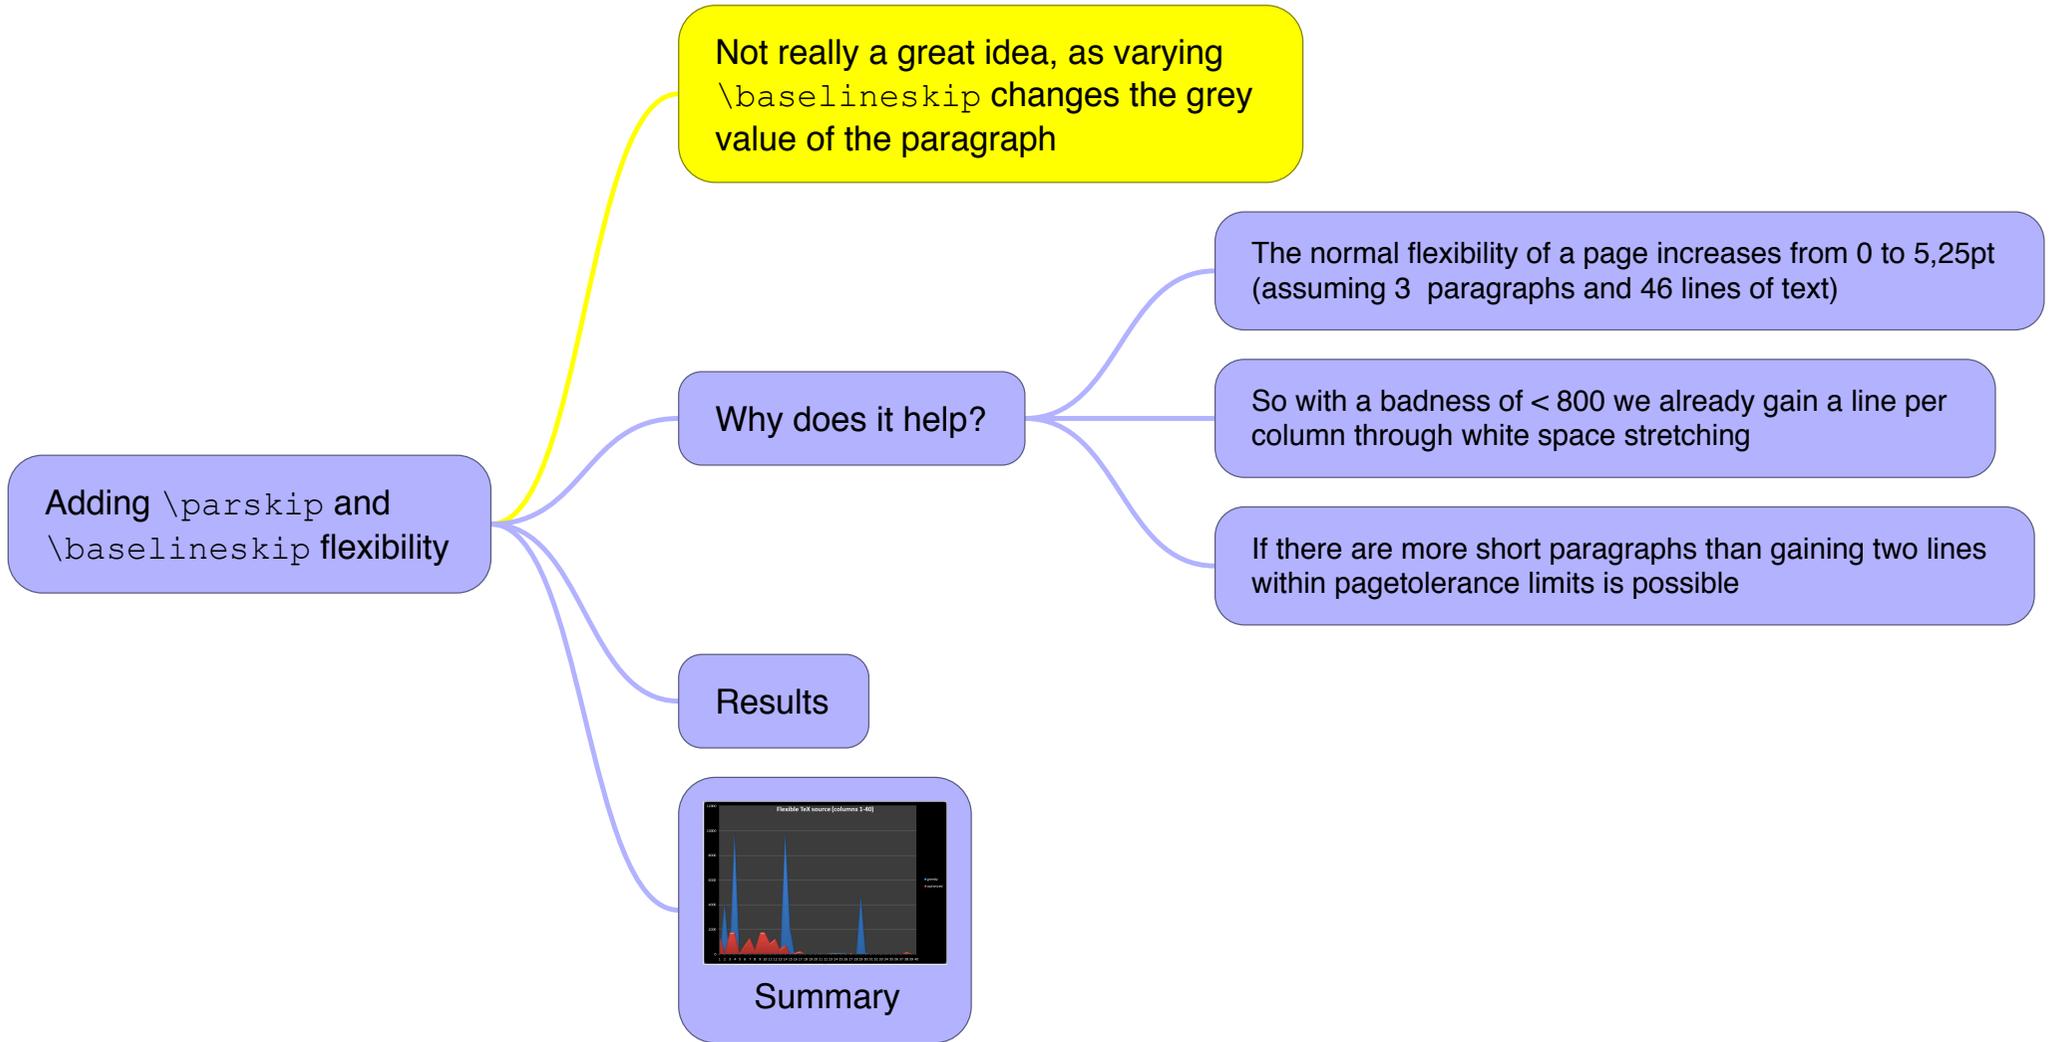
It is also an invalid  
approach if you intend to  
do grid-based typesetting

Why does it help?

Results



Summary



Not really a great idea, as varying `\baselineskip` changes the grey value of the paragraph

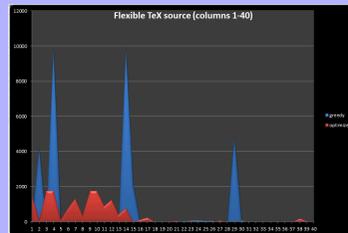
Why does it help?

with standard TeX

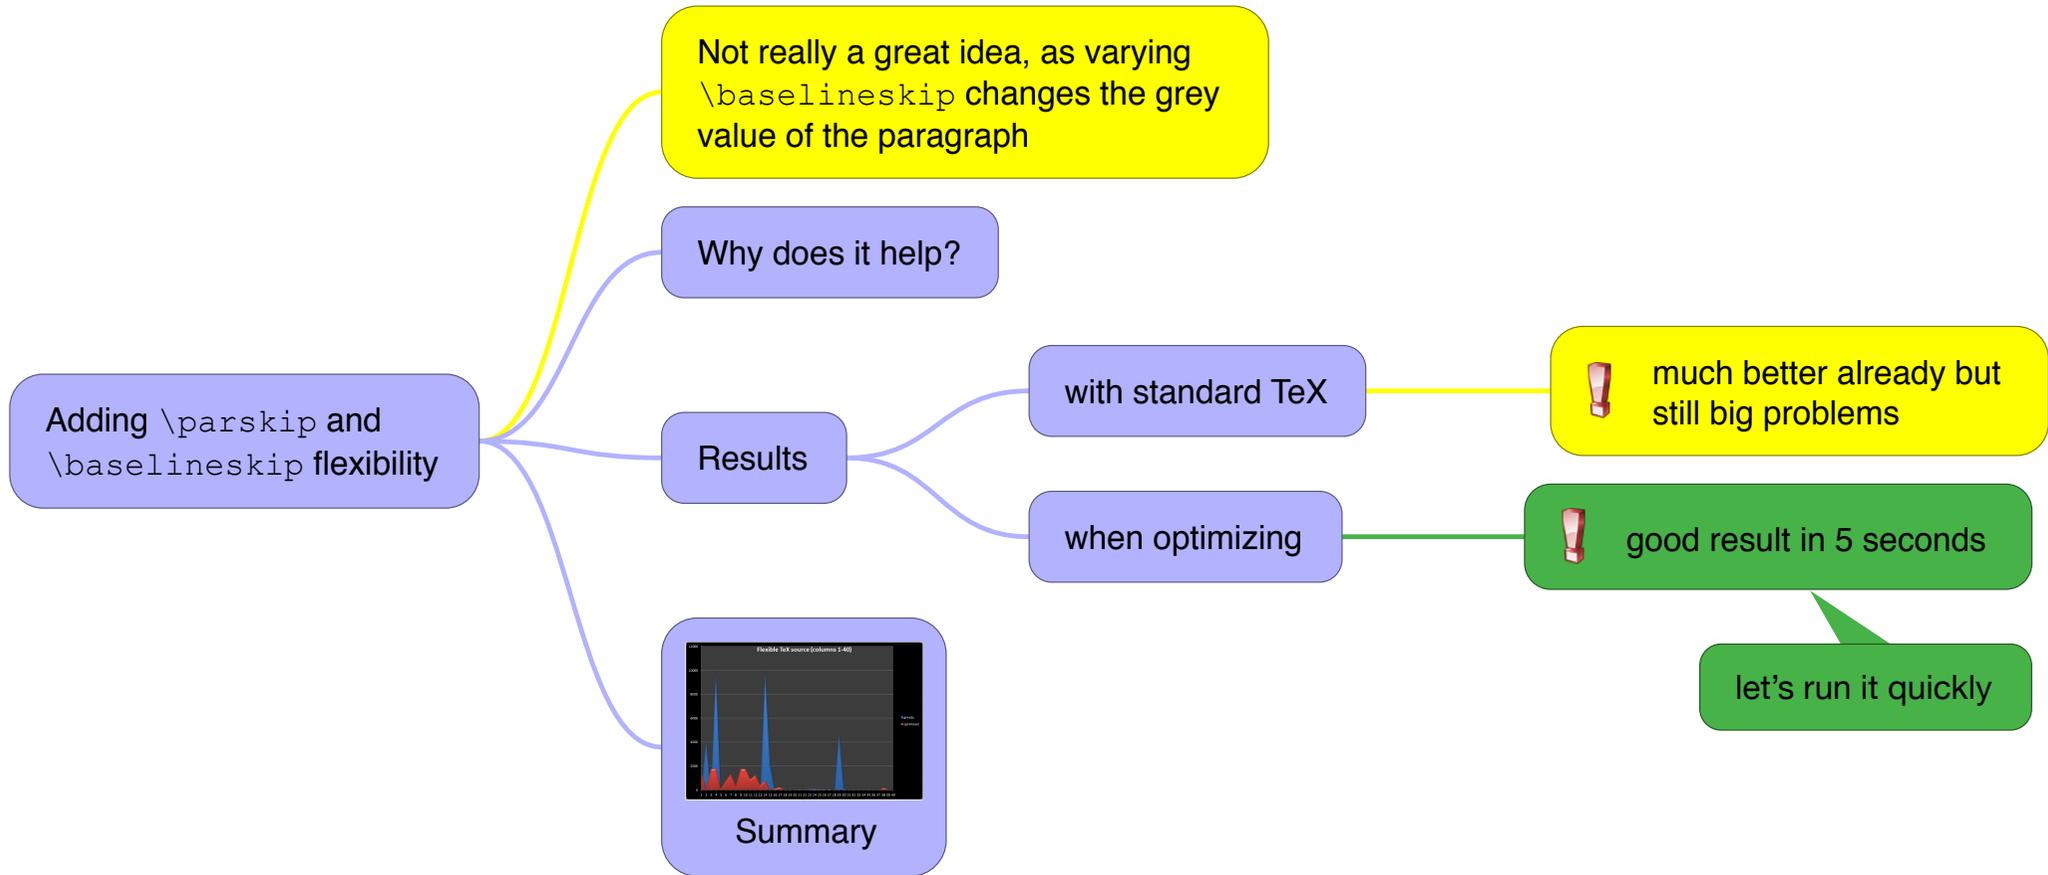
when optimizing

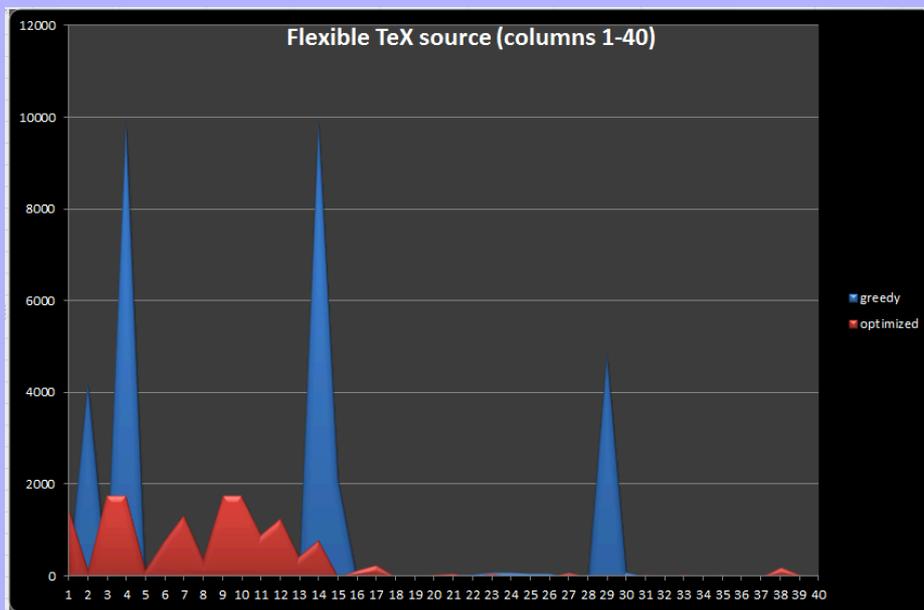
Results

Adding `\parskip` and `\baselineskip` flexibility

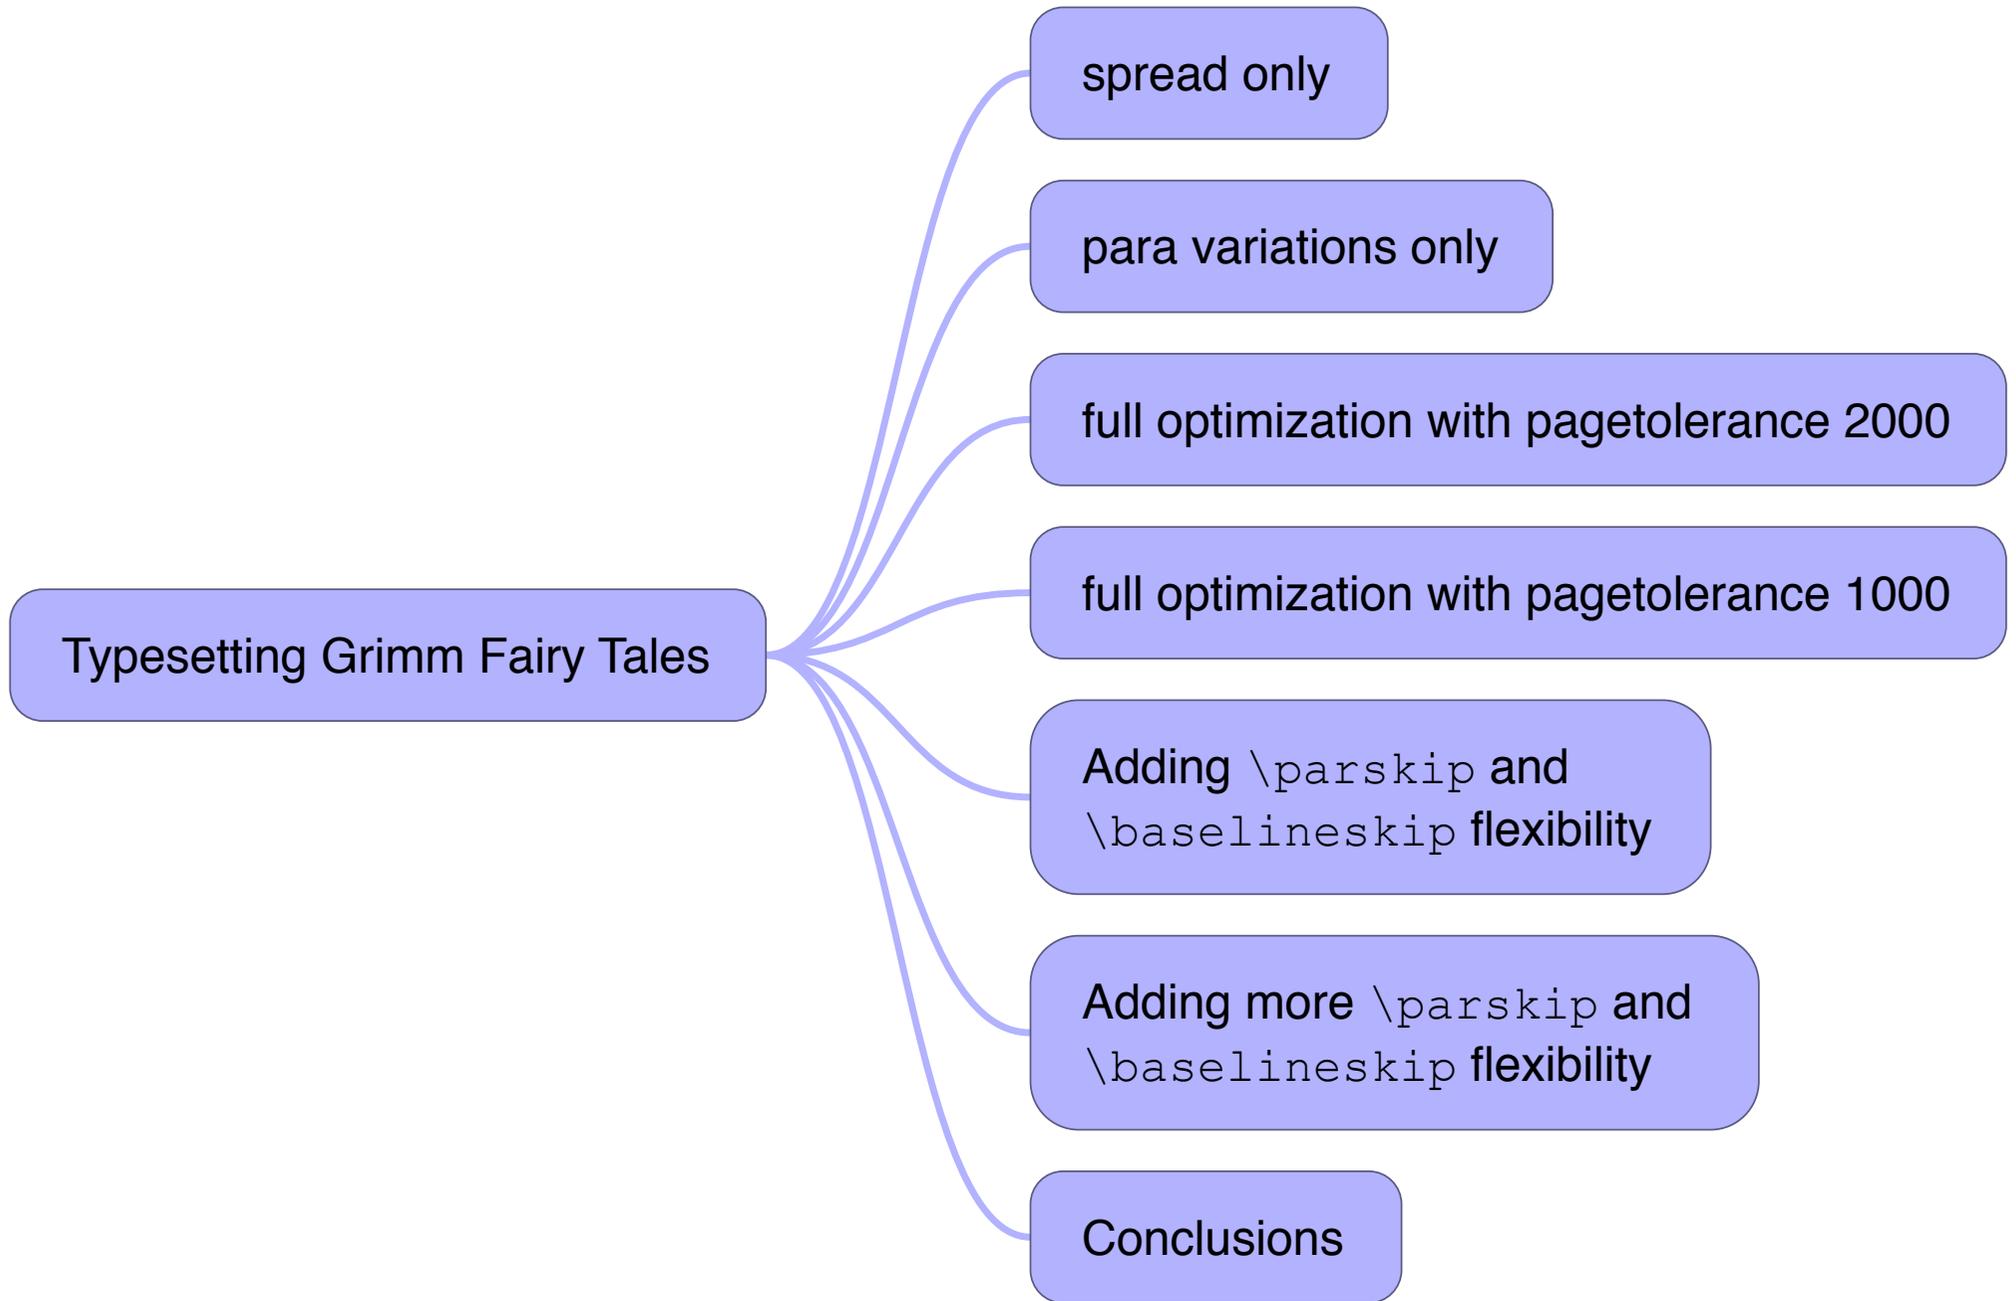


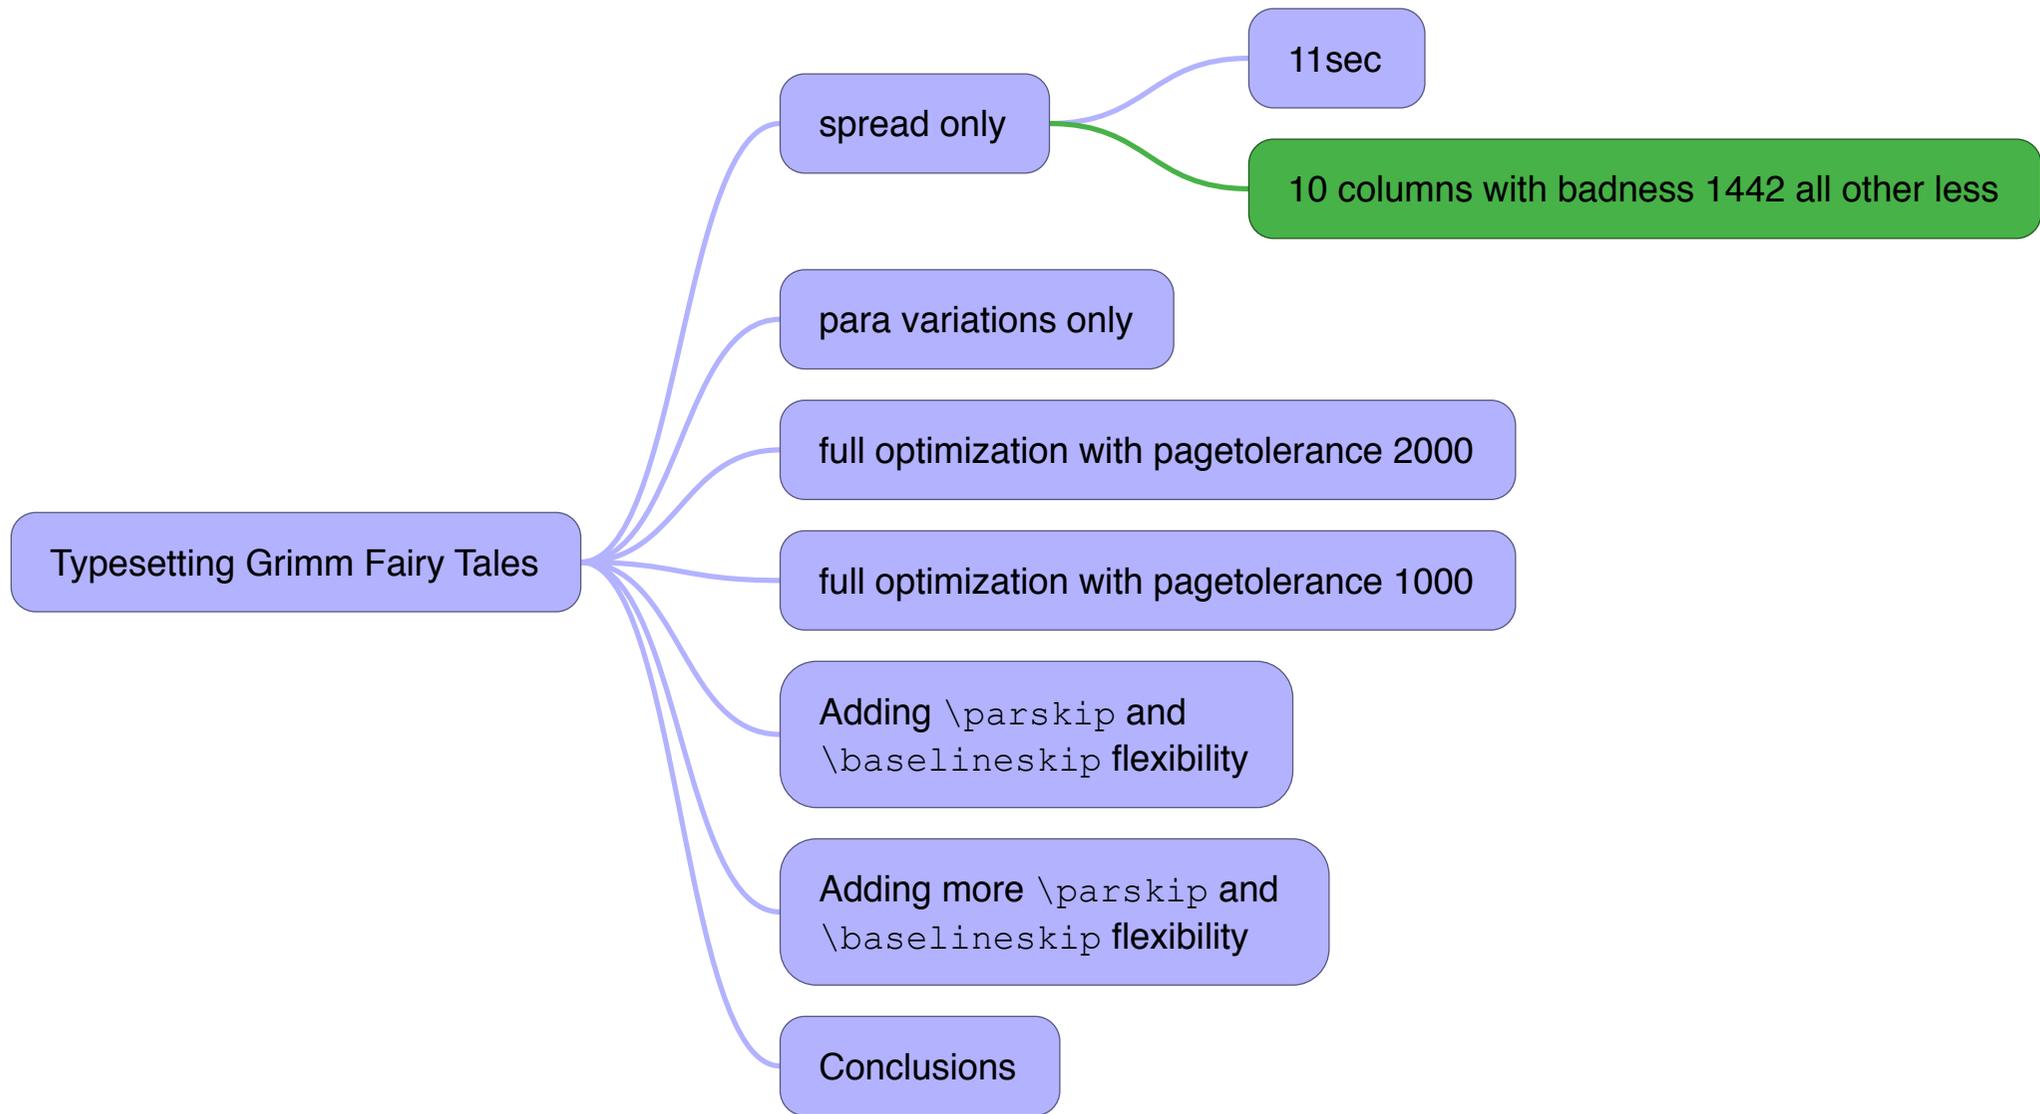
Summary



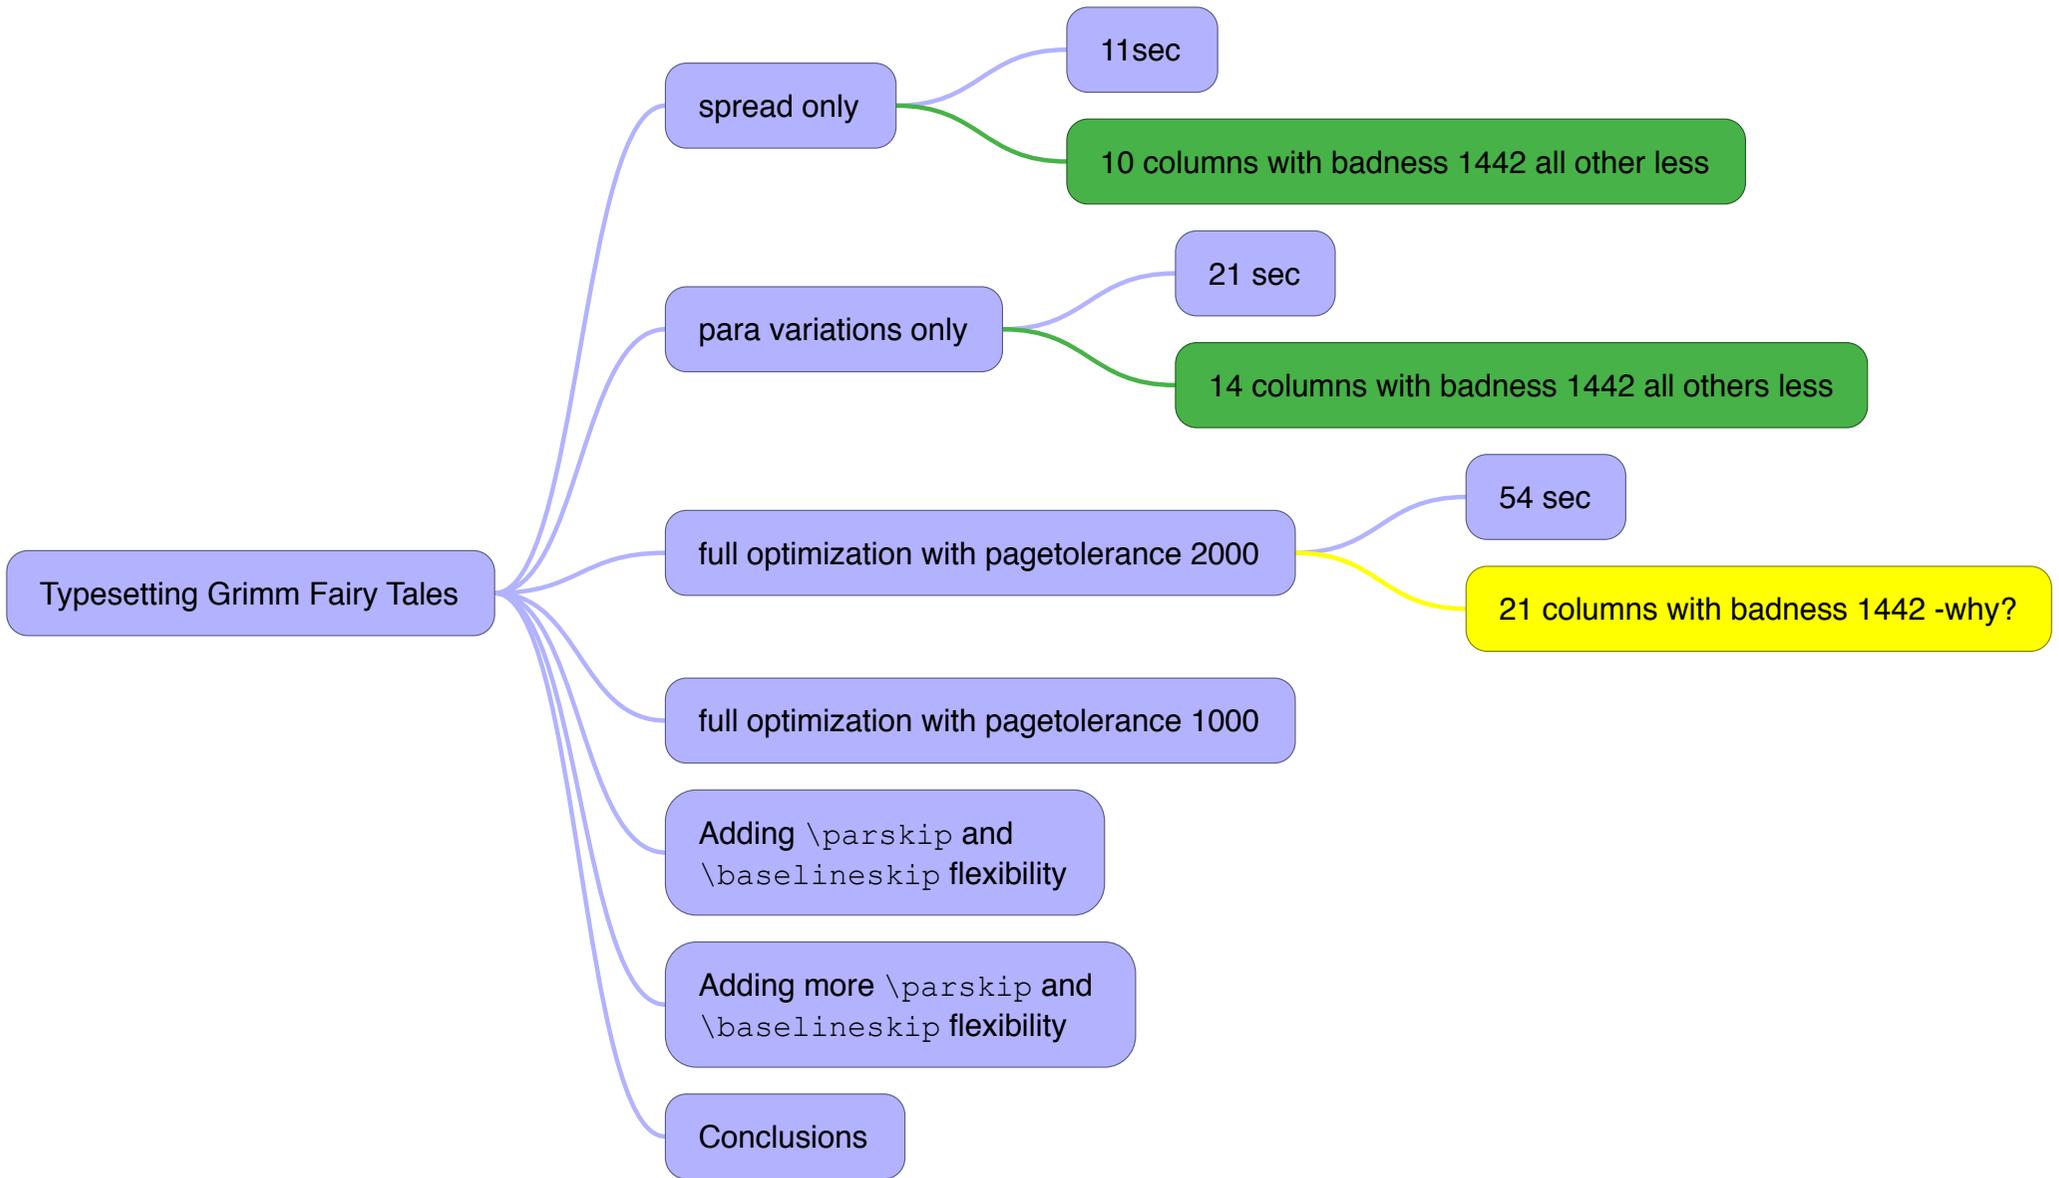


# Summary

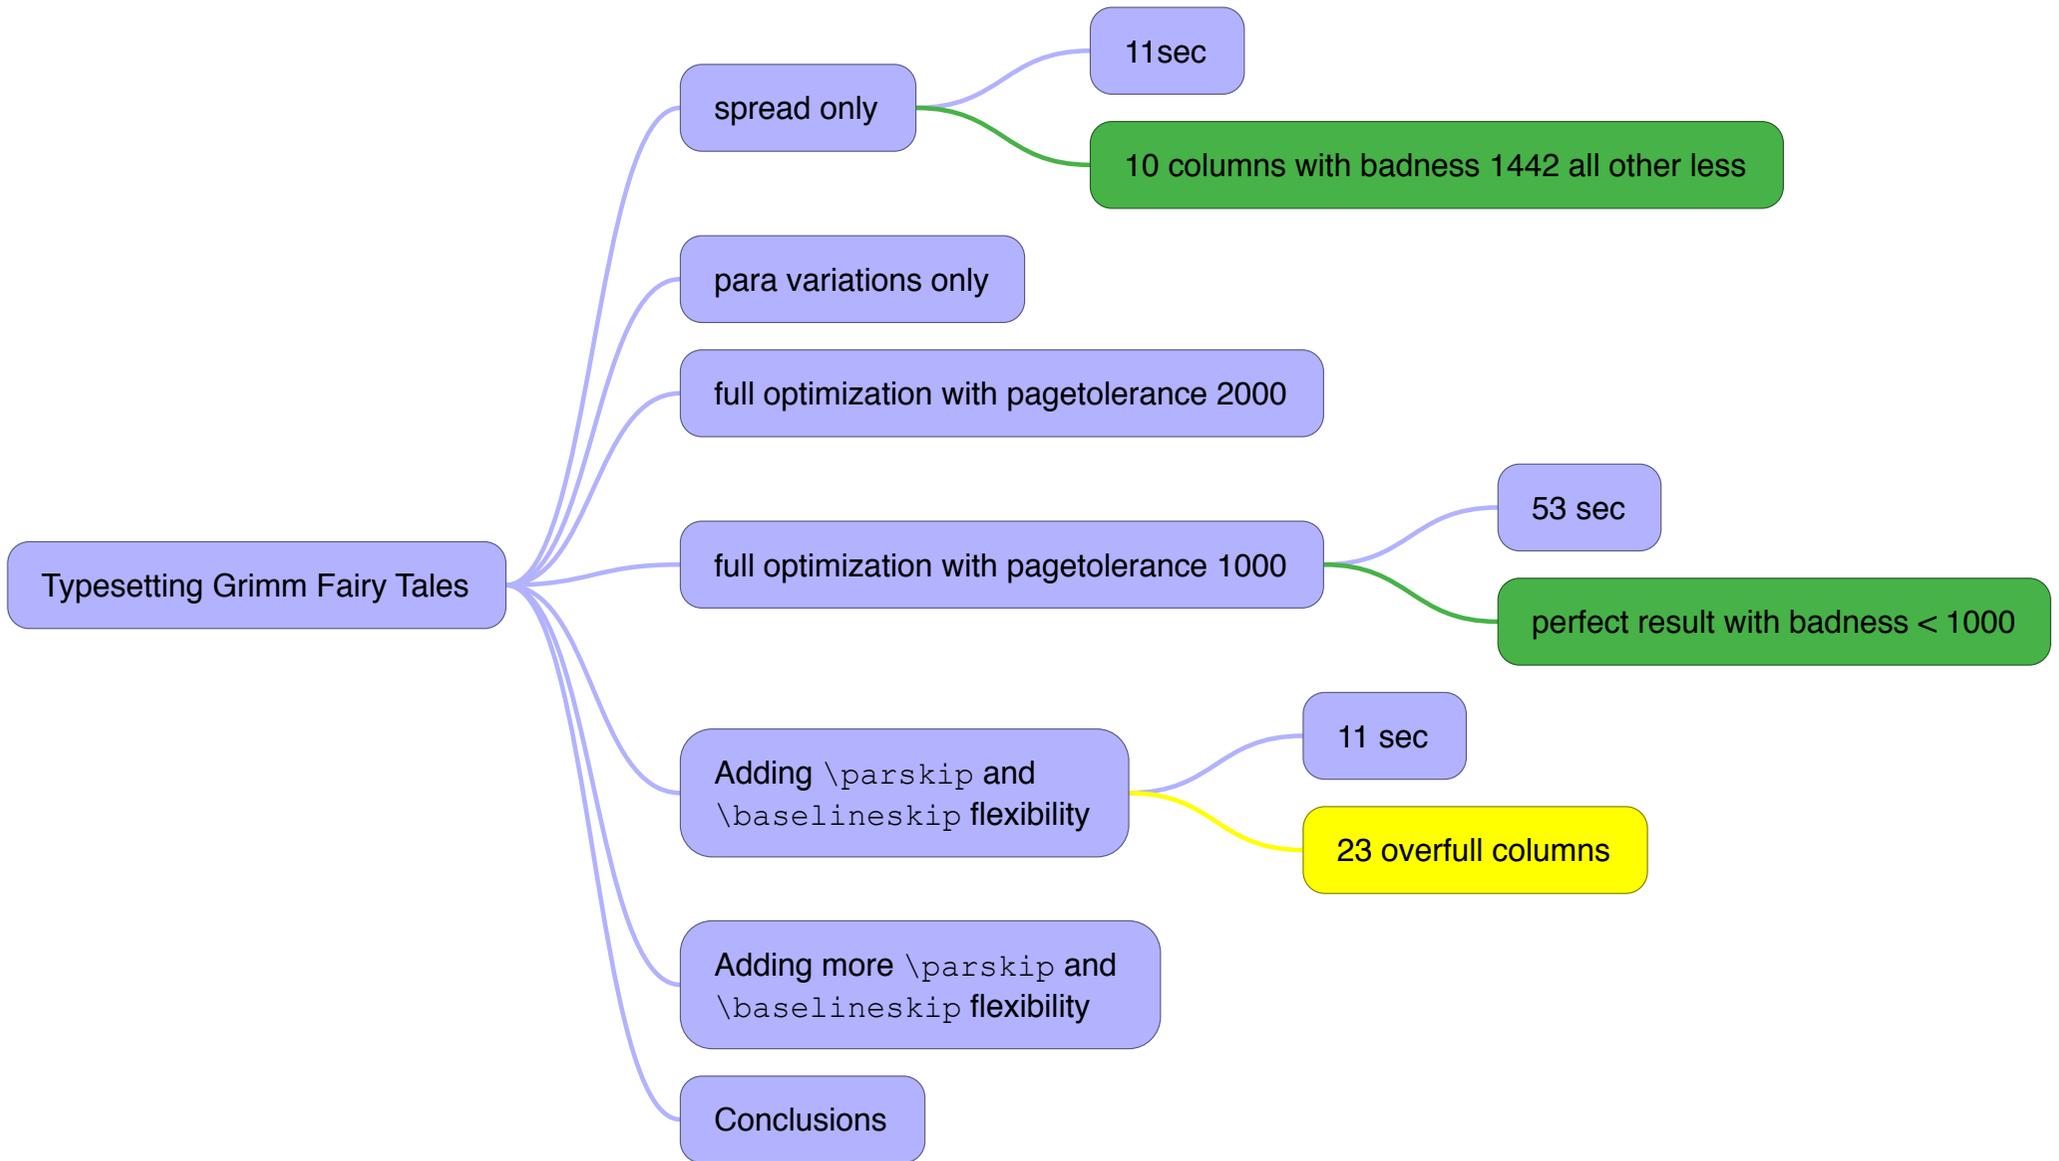






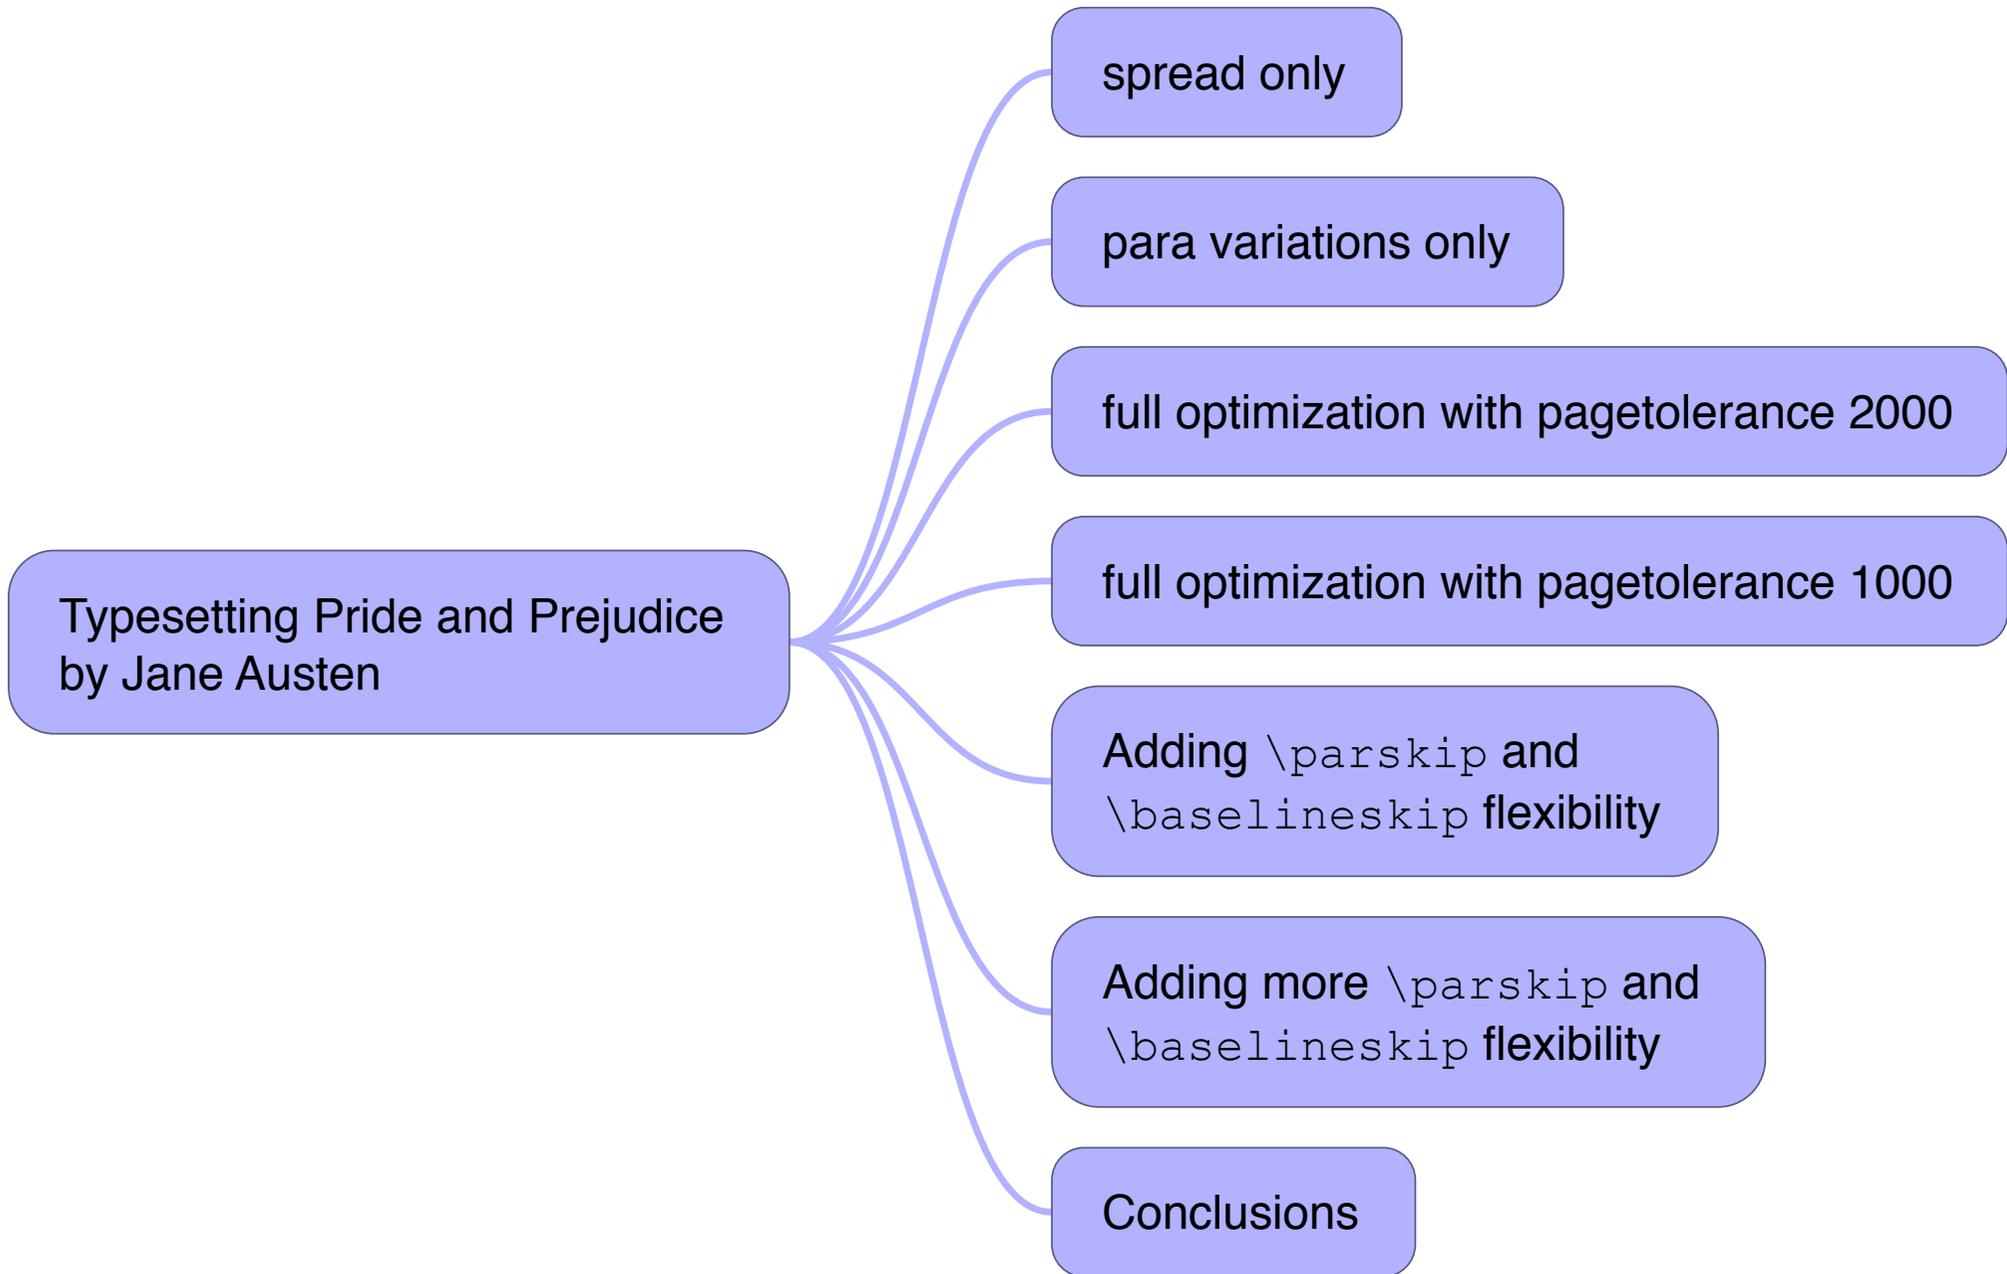


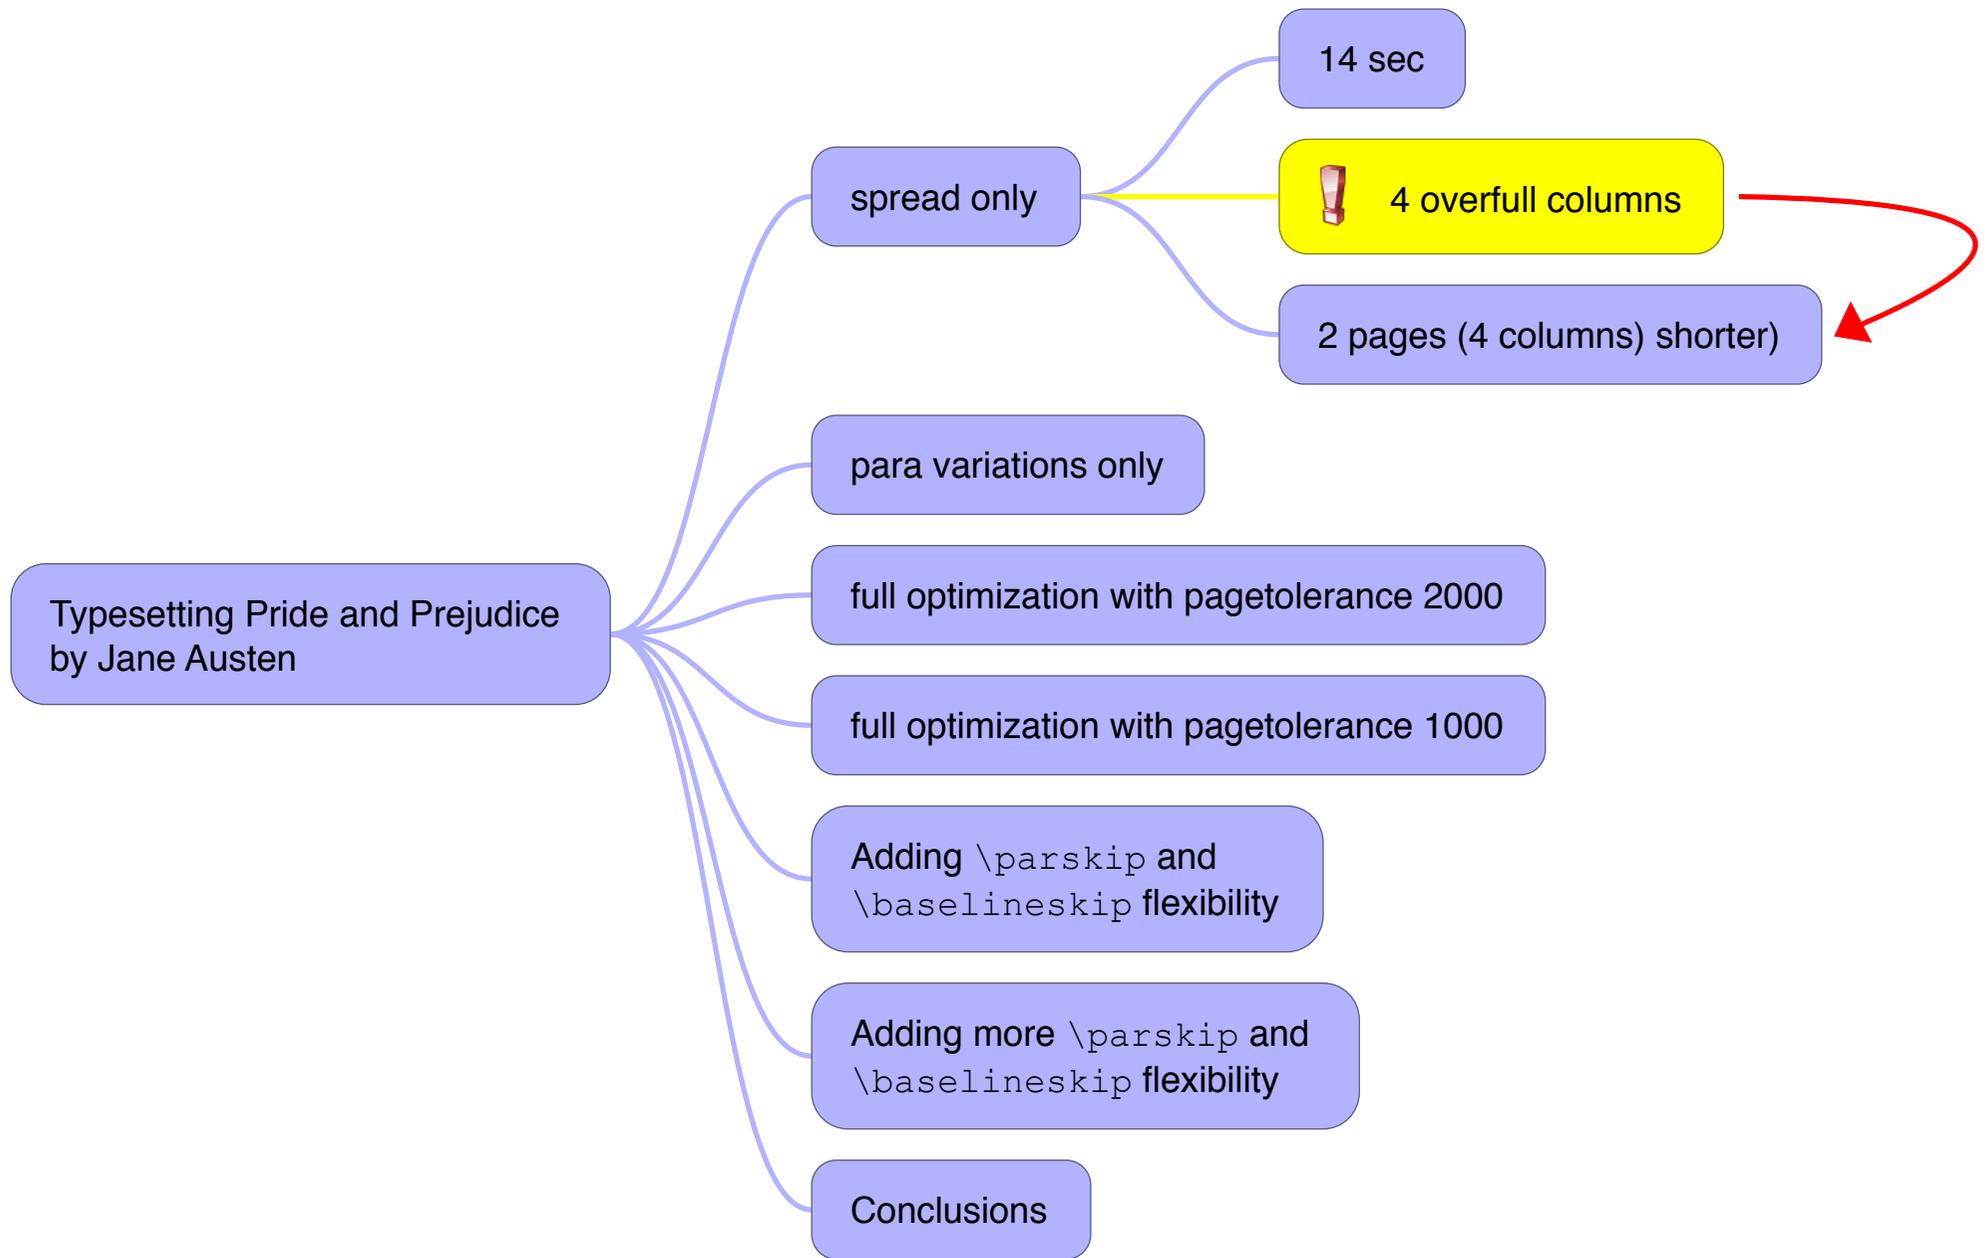


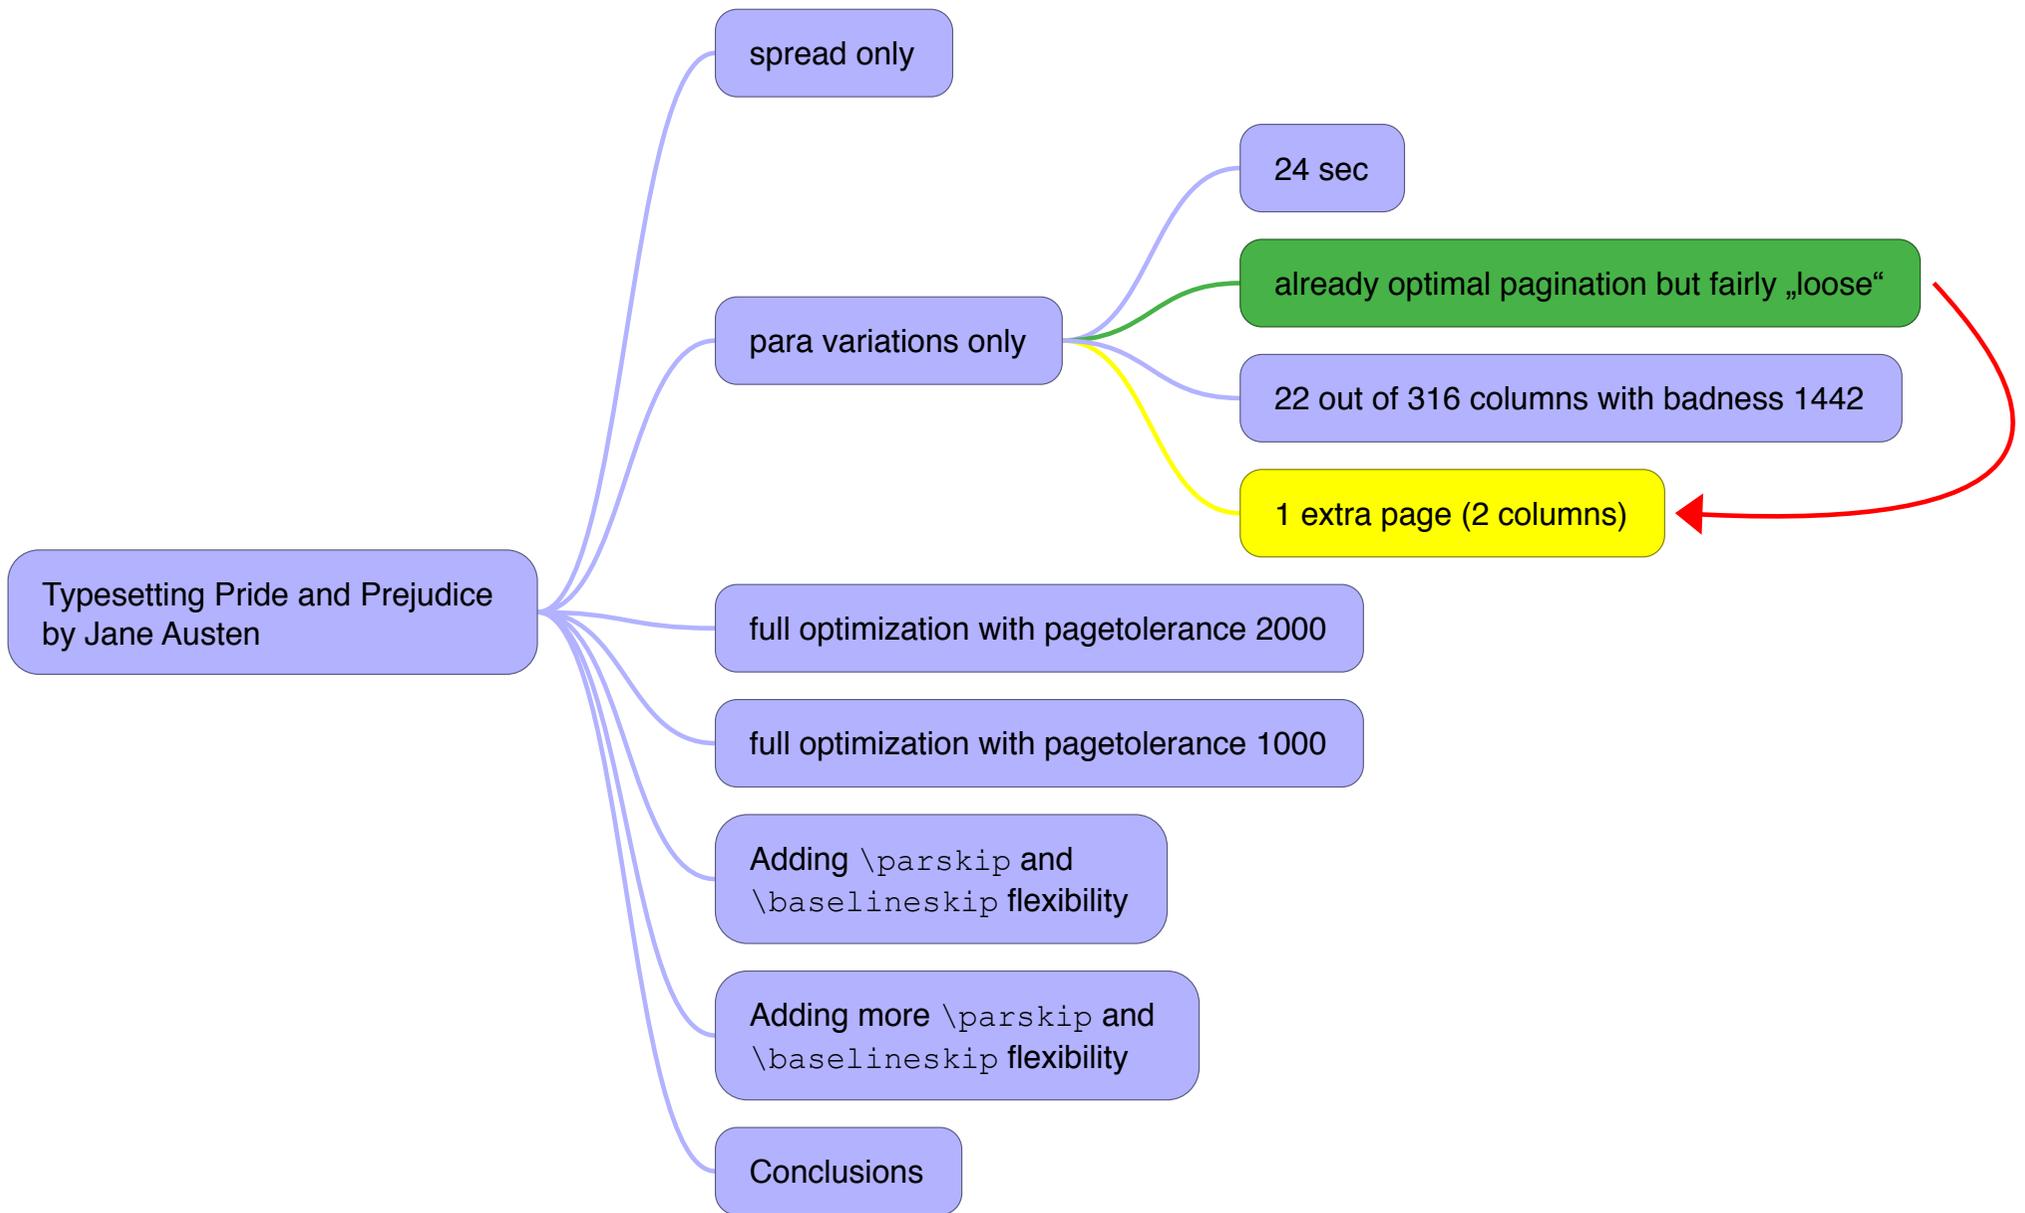


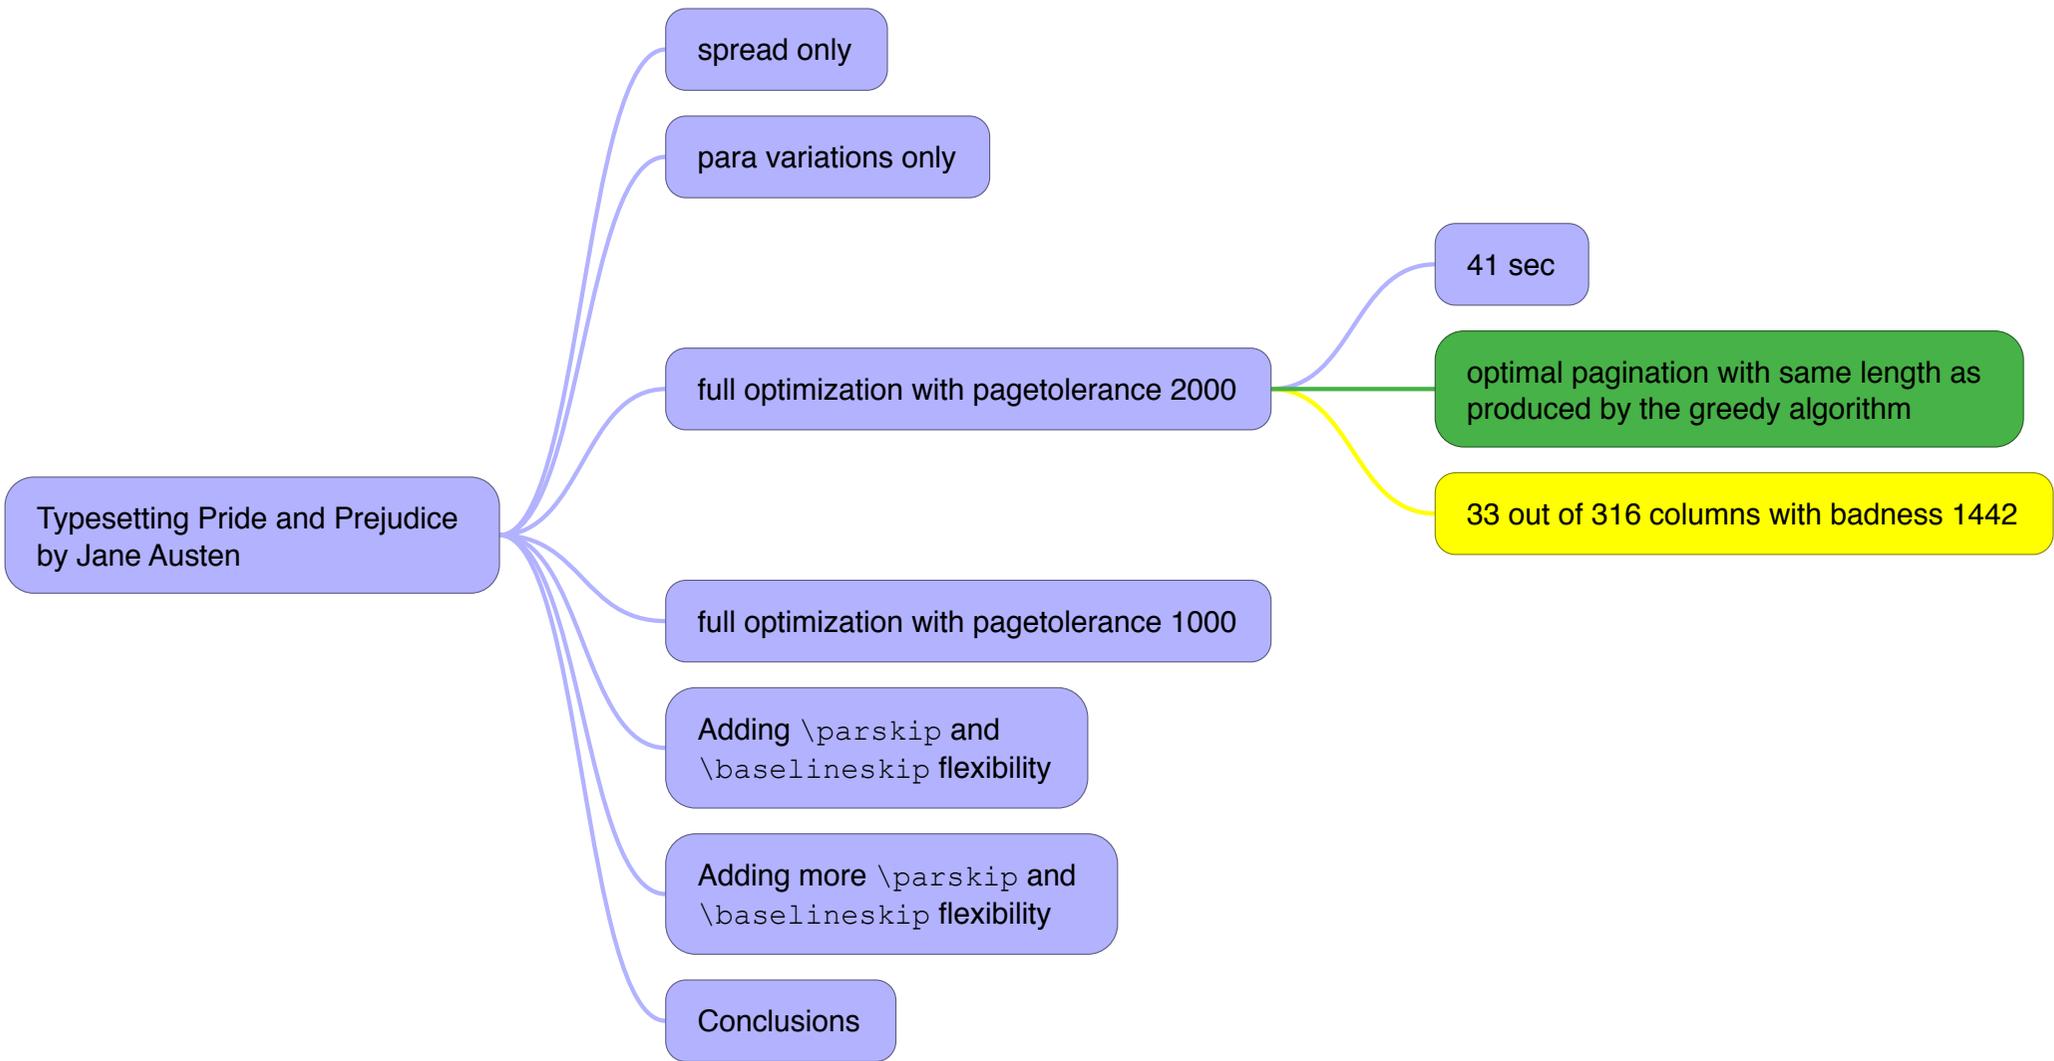


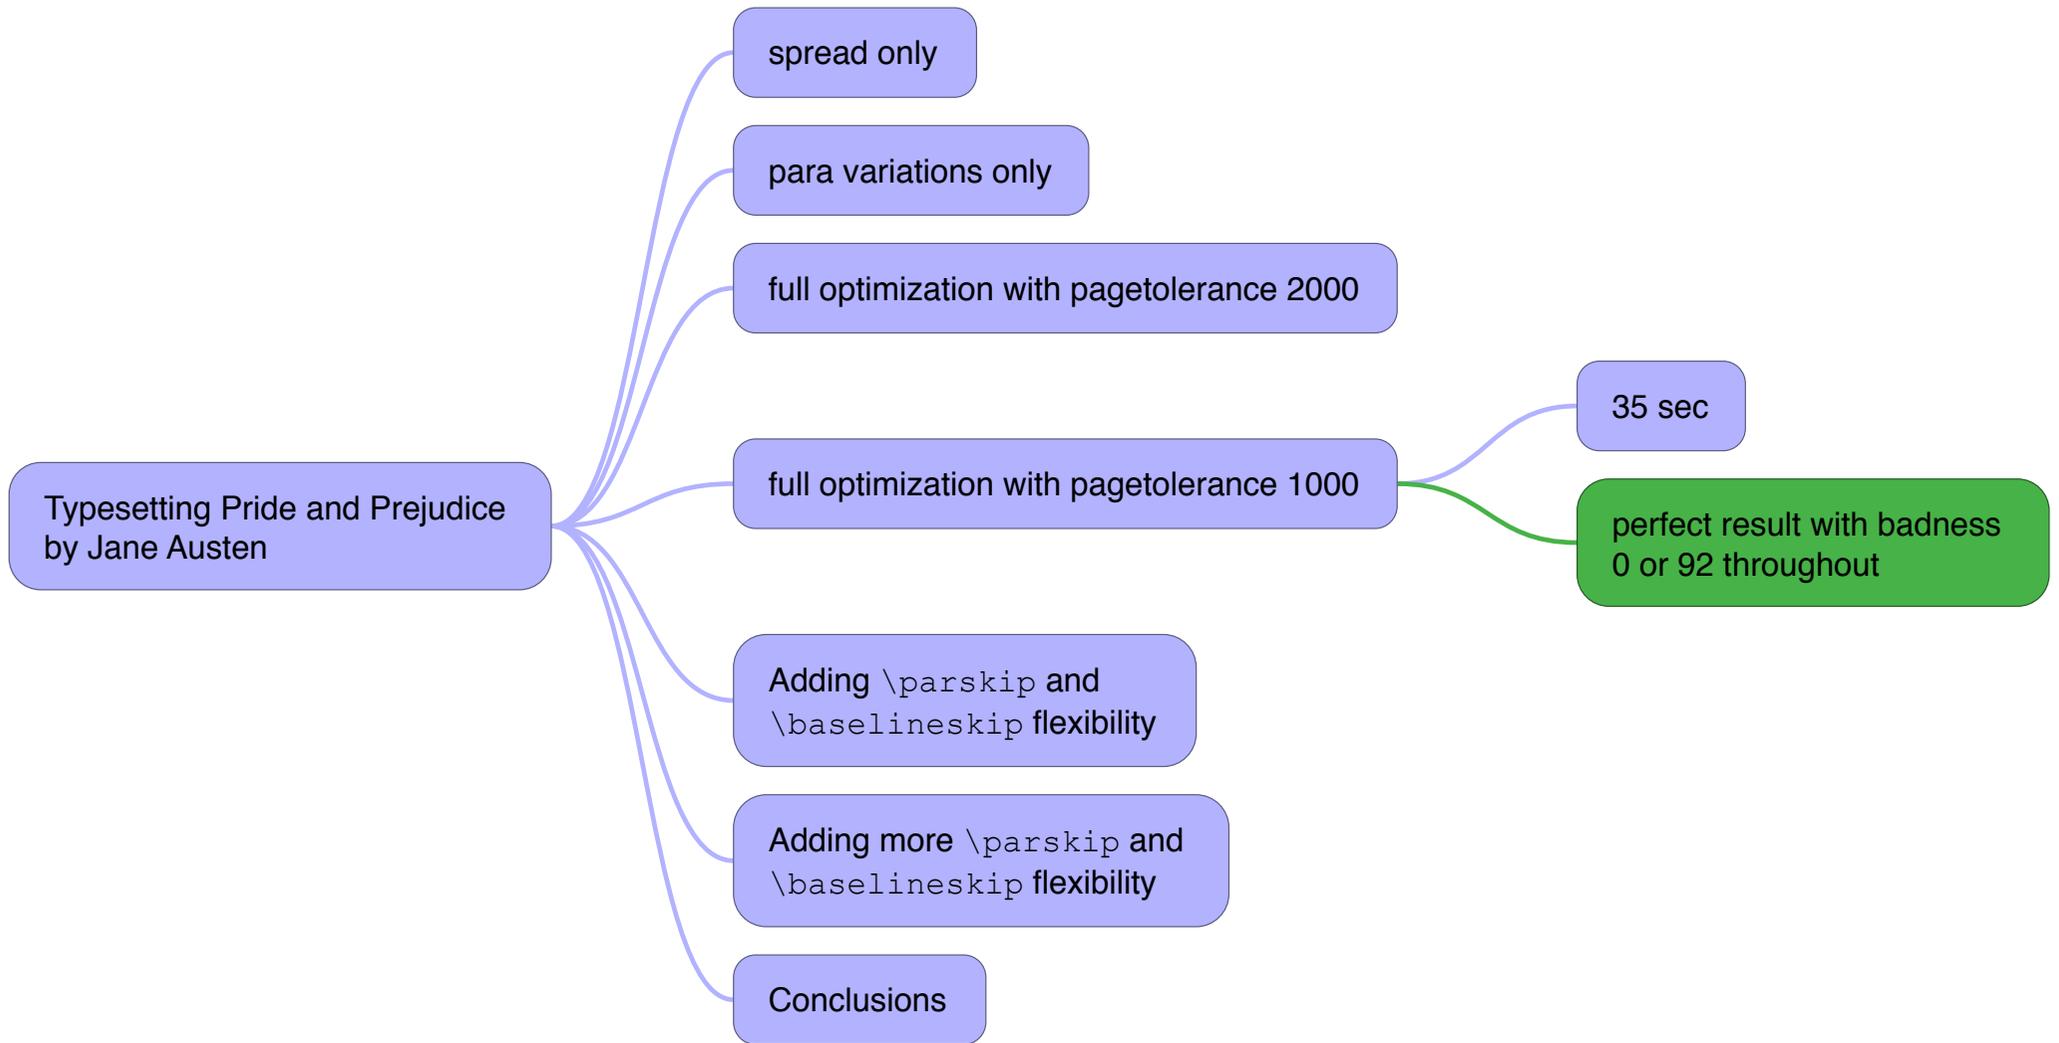












Typesetting Pride and Prejudice by Jane Austen

spread only

para variations only

full optimization with pagetolerance 2000

full optimization with pagetolerance 1000

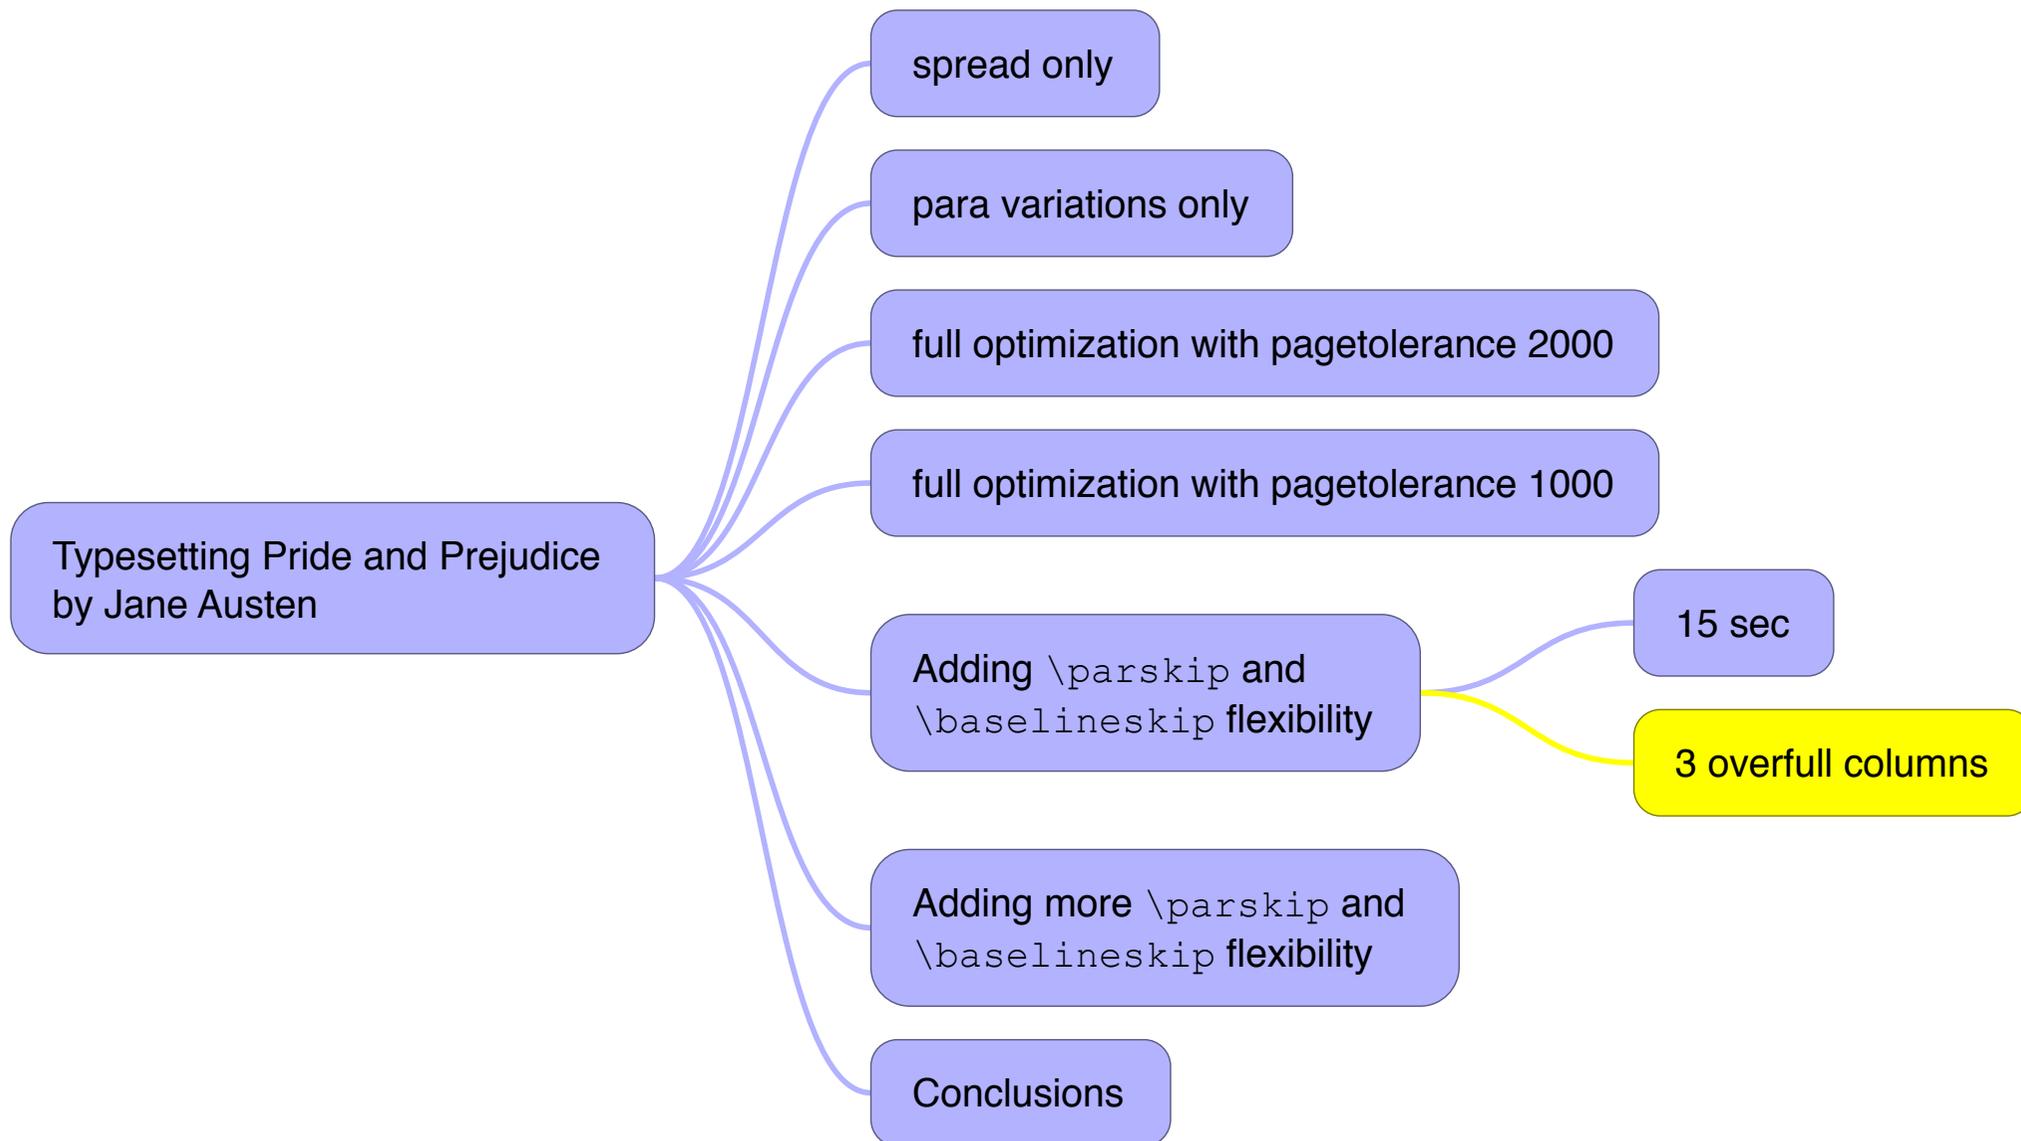
Adding `\parskip` and `\baselineskip` flexibility

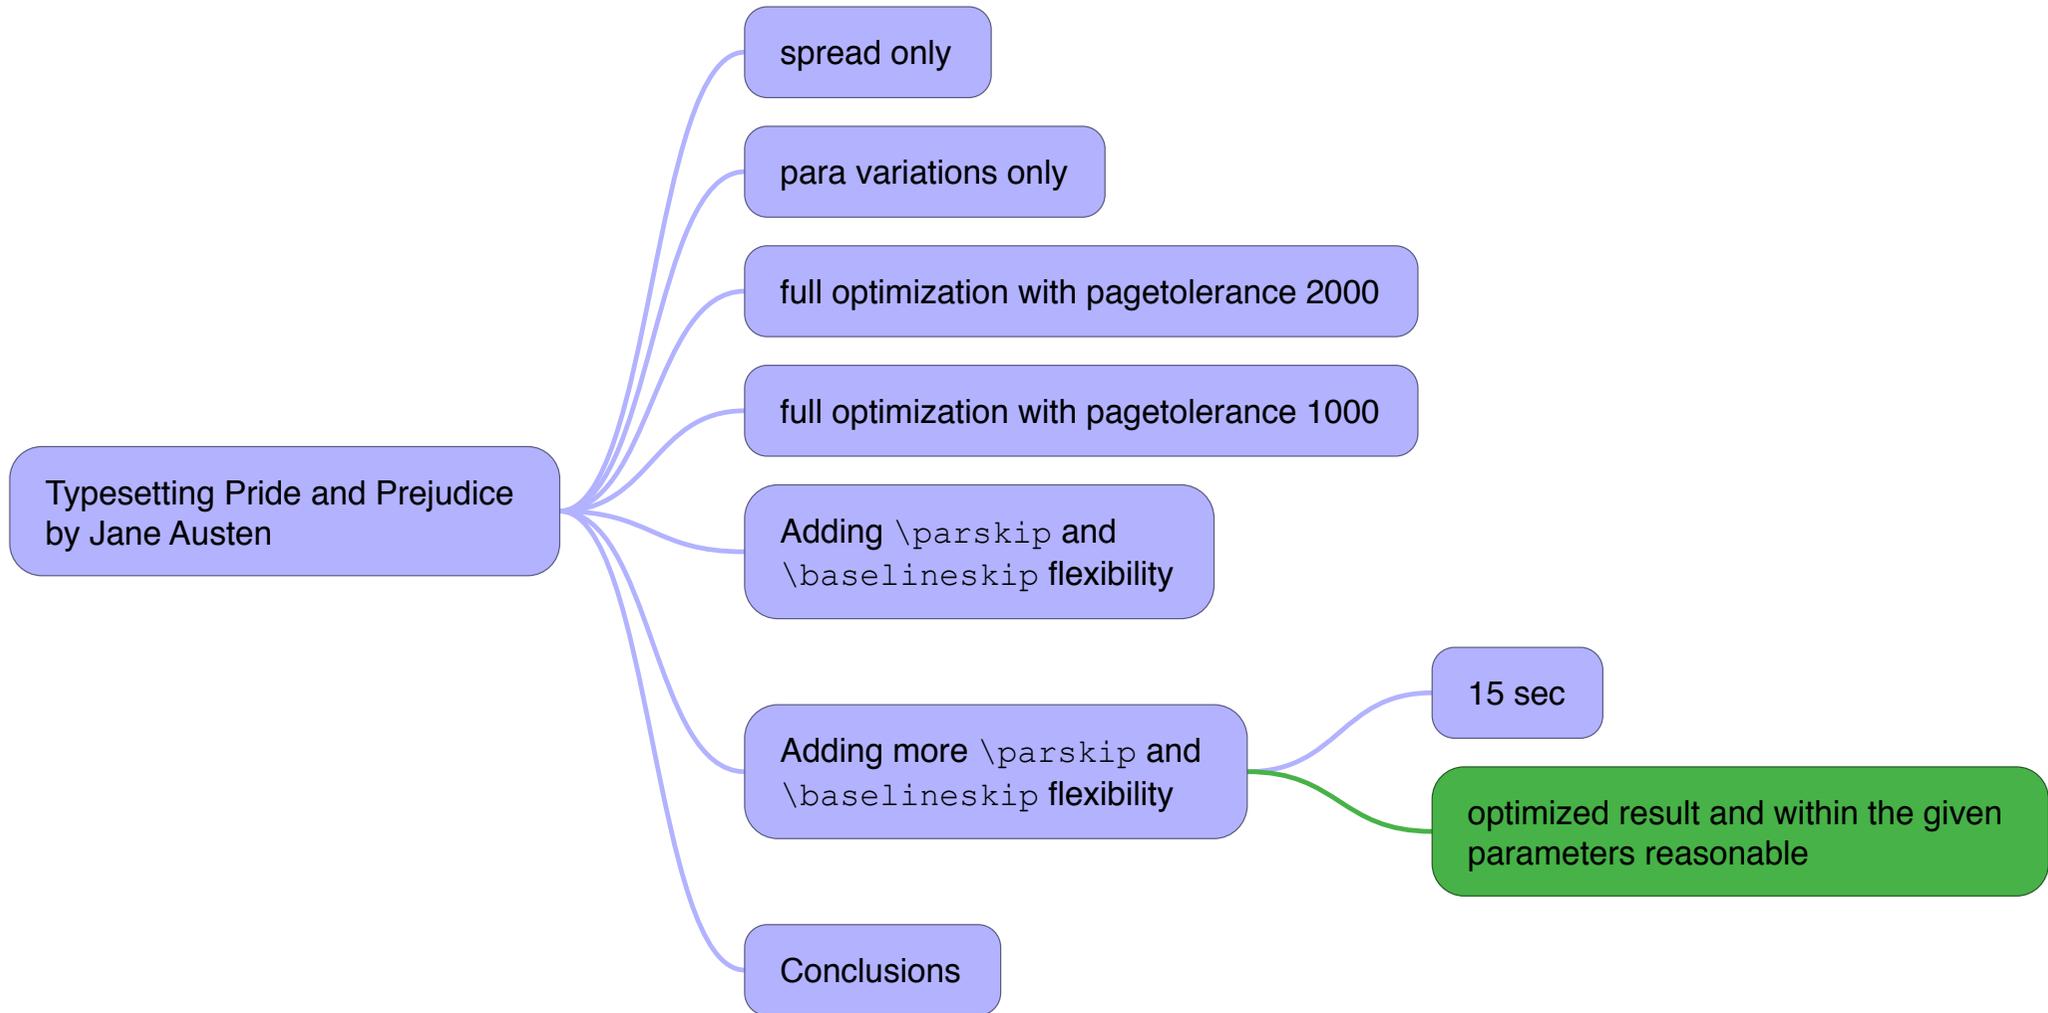
Adding more `\parskip` and `\baselineskip` flexibility

Conclusions

35 sec

perfect result with badness 0 or 92 throughout





Typesetting Pride and Prejudice by Jane Austen

spread only

para variations only

full optimization with pagetolerance 2000

full optimization with pagetolerance 1000

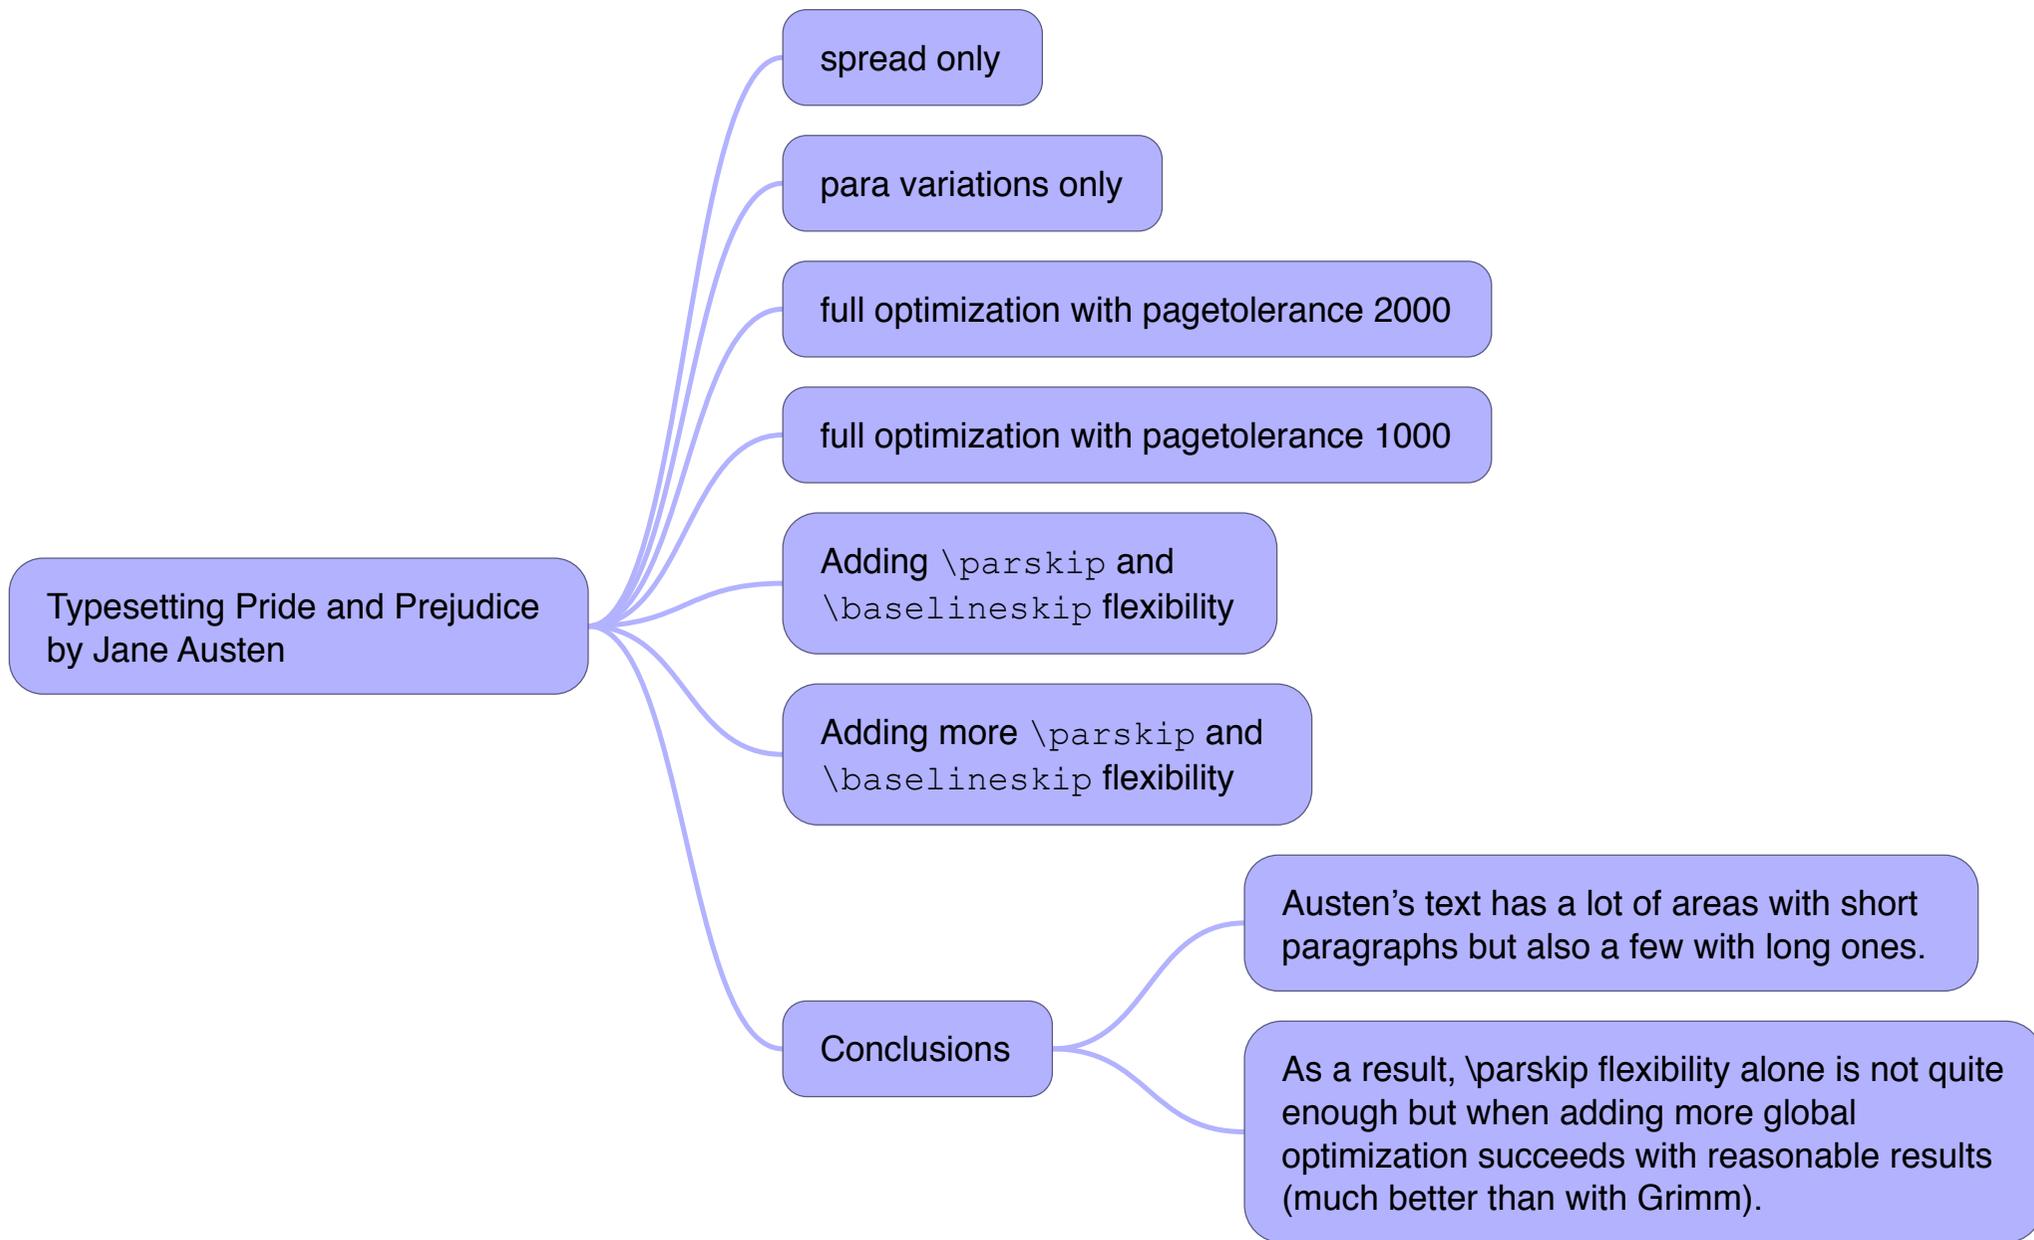
Adding `\parskip` and `\baselineskip` flexibility

Adding more `\parskip` and `\baselineskip` flexibility

Conclusions

15 sec

optimized result and within the given parameters reasonable





The tale of the  
tail of bugs

\looseness **woes**

paragraph variation  
woes

other surprises



The tale of the  
tail of bugs

By default setting `\looseness > 0` gives you very short last lines with a single word in it 😞

Environments like `center` will generate empty extra lines in response to `\looseness > 0` 😞

Setting `\looseness=-1` may result in enlarging the paragraph and vice versa 😞

`\looseness` woes

partly addressed by changing the paragraph material prior to the `\looseness` trials so that the last set of word spaces get additional penalties attached



partly because TeX may still decide to hyphenate the last word and make a very short last line (`\finalhyphendemerits` helps but ...)



By default setting `\looseness > 0` gives you very short last lines with a single word in it



Environments like `center` will generate empty extra lines in response to `\looseness > 0`



Setting `\looseness=-1` may result in enlarging the paragraph and vice versa



`\looseness` woes

this needs addressing at the macro code level by providing a flag that tells phase 2 not to apply `\looseness` changes to a certain paragraph



By default setting `\looseness > 0` gives you very short last lines with a single word in it



Environments like `center` will generate empty extra lines in response to `\looseness > 0`



Setting `\looseness=-1` may result in enlarging the paragraph and vice versa



`\looseness` woes

By default setting `\looseness > 0` gives you very short last lines with a single word in it 😞

Environments like `center` will generate empty extra lines in response to `\looseness > 0` 😞

Setting `\looseness=-1` may result in enlarging the paragraph and vice versa 😞

`\looseness` woes

addressed by carefully checking the results and not just assuming that the TeX reporting the `\looseness` setting worked actually means that it work as expected ✓

Different paragraph variations may result in different values of `\prevdepth` after the paragraph, but the material following is always set as if `\looseness=0` was used



paragraph variation  
woes

(Lua)TeX may report success setting a paragraph with `\looseness=-1` even though it is invalid as it contains an overfull line



addressed by monitoring all variation for a change in `\prevdepth` and adjusting their output, so that they all can continue in the same way



Different paragraph variations may result in different values of `\prevdepth` after the paragraph, but the material following is always set as if `\looseness=0` was used



(Lua)TeX may report success setting a paragraph with `\looseness=-1` even though it is invalid as it contains an overfull line



paragraph variation woes

This will happen if already the `\looseness=0` case is invalid. As the algorithm supports different values for `\tolerance` (without and with `\looseness`) this is a possibility.



addressed by checking for it



Different paragraph variations may result in different values of `\prevdepth` after the paragraph, but the material following is always set as if `\looseness=0` was used



(Lua)TeX may report success setting a paragraph with `\looseness=-1` even though it is invalid as it contains an overfull line



paragraph variation woes

The greedy algorithm only produces underfull pages, the optimizing algorithm can **and will** produce overfulls



`\clearpage` invokes the output routine



It is very difficult if not impossible not to introduce some inaccuracies due to rounding differences between the TeX engine and external processing in Lua



In a number of cases when I was searching for bugs in my code I found them eventually in ...



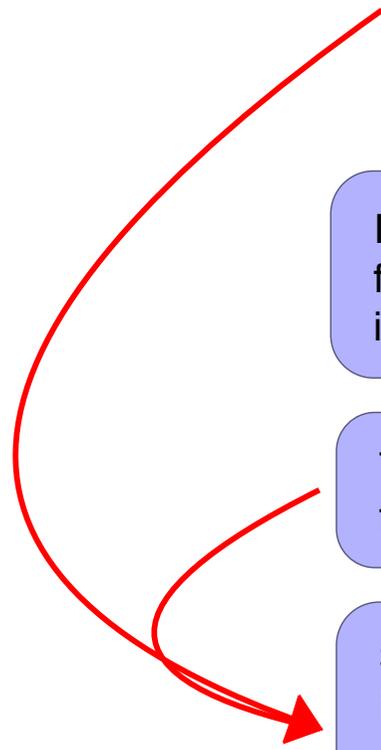
The final results were often quite different from the expectation



**Special case:** If the flexibility on the page is minimal and the page height is not exactly an integral number of lines then the results get very strange



other surprises



perhaps one should add an option to use the best infeasible break instead 

as a minimum some sort of `\overfullrule` indicator should be implemented 

The greedy algorithm only produces underfull pages, the optimizing algorithm can **and will** produce overfulls 

`\clearpage` invokes the output routine 

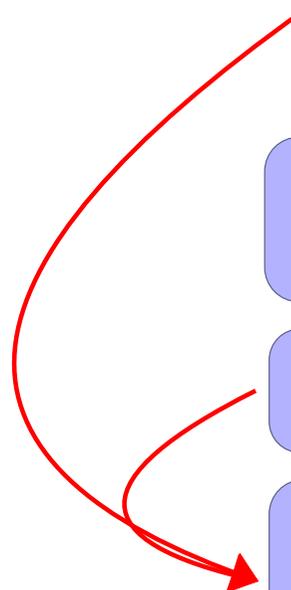
It is very difficult if not impossible not to introduce some inaccuracies due to rounding differences between the TeX engine and external processing in Lua 

In a number of cases when I was searching for bugs in my code I found them eventually in ...   


The final results were often quite different from the expectation 

**Special case:** If the flexibility on the page is minimal and the page height is not exactly an integral number of lines then the results get very strange 

other surprises



floats too :-) This needs addressing in the next round of implementation



The greedy algorithm only produces underfull pages, the optimizing algorithm can **and will** produce overfulls



`\clearpage` invokes the output routine



It is very difficult if not impossible not to introduce some inaccuracies due to rounding differences between the TeX engine and external processing in Lua



In a number of cases when I was searching for bugs in my code I found them eventually in ...



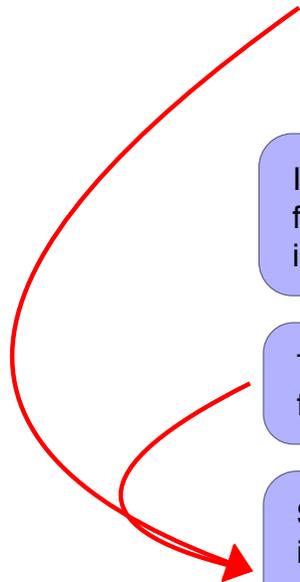
The final results were often quite different from the expectation



**Special case:** If the flexibility on the page is minimal and the page height is not exactly an integral number of lines then the results get very strange



other surprises





the LuaTeX callback interfaces 😊

The greedy algorithm only produces underfull pages, the optimizing algorithm can **and will** produce overfulls ❓

`\clearpage` invokes the output routine 😞

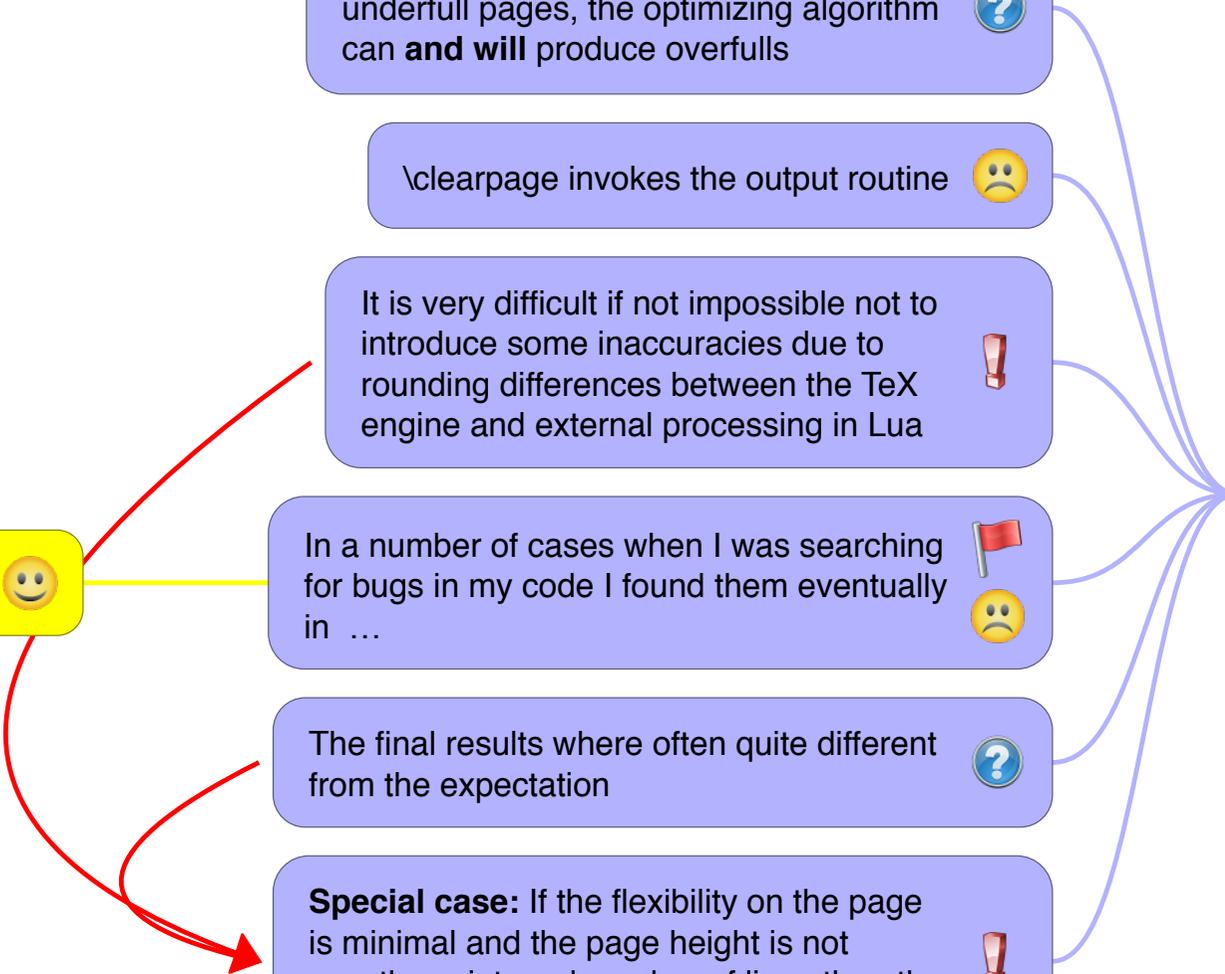
It is very difficult if not impossible not to introduce some inaccuracies due to rounding differences between the TeX engine and external processing in Lua ⚠️

In a number of cases when I was searching for bugs in my code I found them eventually in ... 🇷🇺 😞

The final results were often quite different from the expectation ❓

**Special case:** If the flexibility on the page is minimal and the page height is not exactly an integral number of lines then the results get very strange ⚠️

other surprises



Basically this means that the interaction between different parameter settings isn't fully understood (yet) 

It also shows that a single badness function can model complex aesthetics only to a certain extent 

The greedy algorithm only produces underfull pages, the optimizing algorithm can **and will** produce overfulls 

`\clearpage` invokes the output routine 

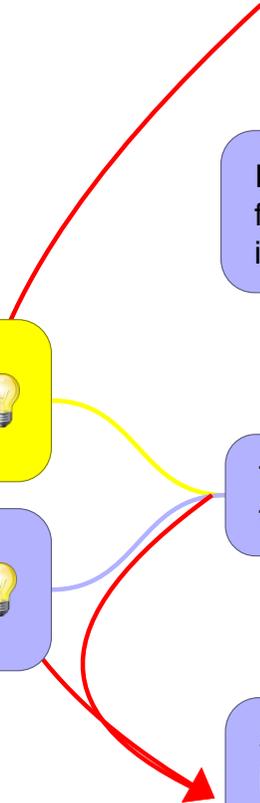
It is very difficult if not impossible not to introduce some inaccuracies due to rounding differences between the TeX engine and external processing in Lua 

In a number of cases when I was searching for bugs in my code I found them eventually in ...   

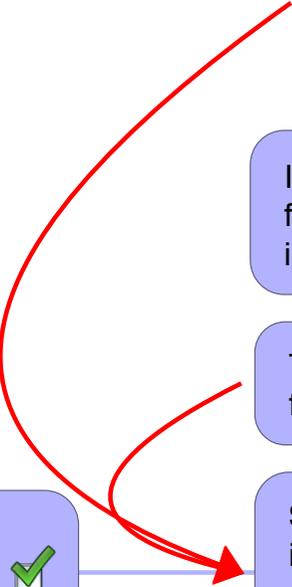

The final results were often quite different from the expectation 

**Special case:** If the flexibility on the page is minimal and the page height is not exactly an integral number of lines then the results get very strange 

other surprises



If one thinks about it is actually quite clear why, but first it got me puzzled a couple of times. 



The greedy algorithm only produces underfull pages, the optimizing algorithm can **and will** produce overfulls 

`\clearpage` invokes the output routine 

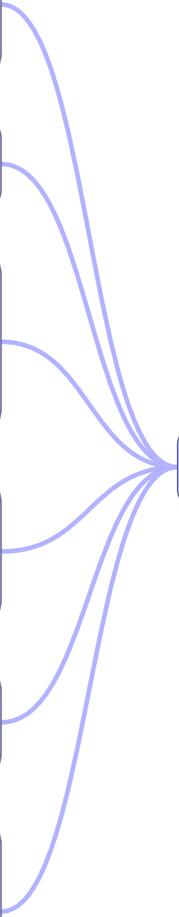
It is very difficult if not impossible not to introduce some inaccuracies due to rounding differences between the TeX engine and external processing in Lua 

In a number of cases when I was searching for bugs in my code I found them eventually in ...   


The final results were often quite different from the expectation 

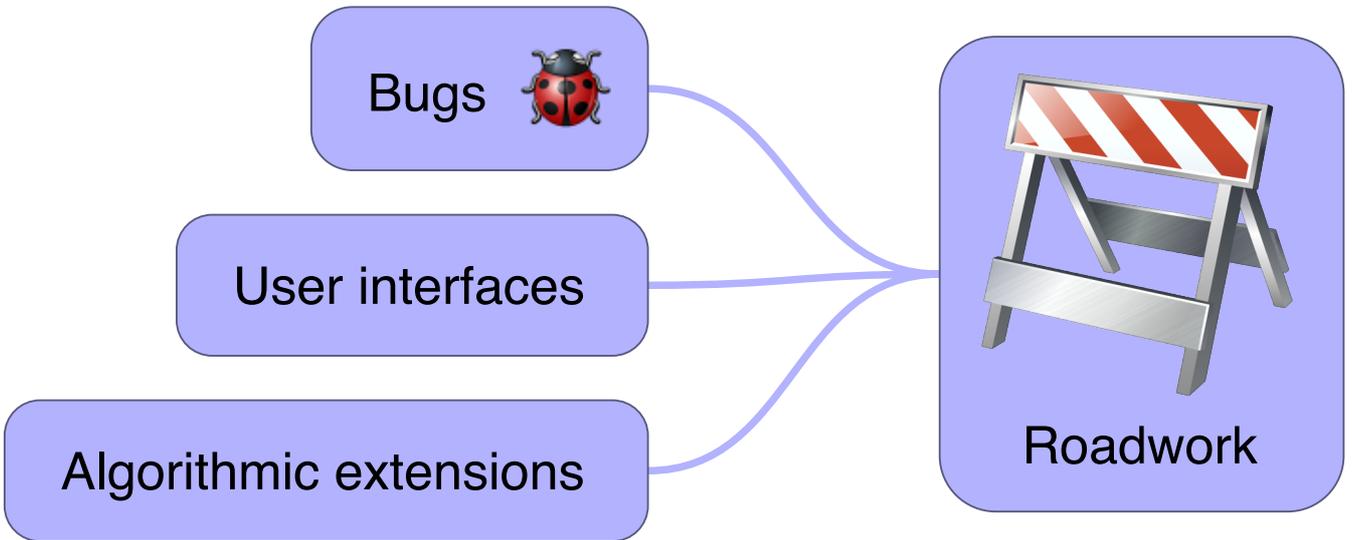
**Special case:** If the flexibility on the page is minimal and the page height is not exactly an integral number of lines then the results get very strange 

other surprises





**Roadwork**



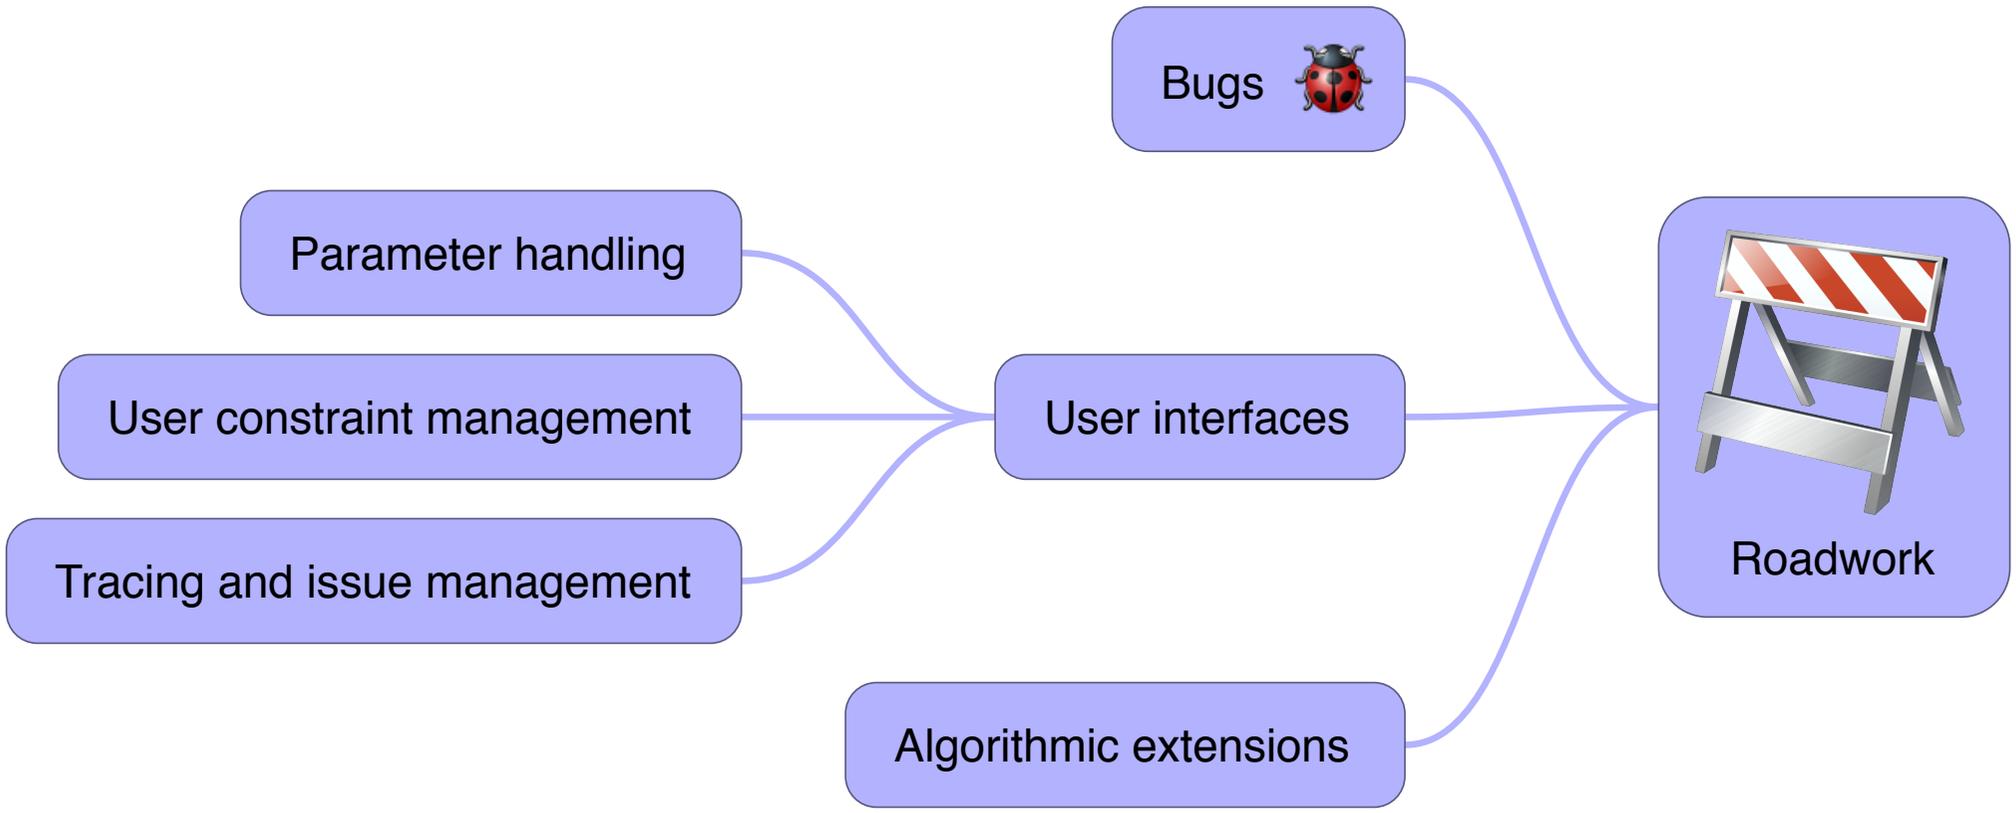
Of course there aren't any :-)

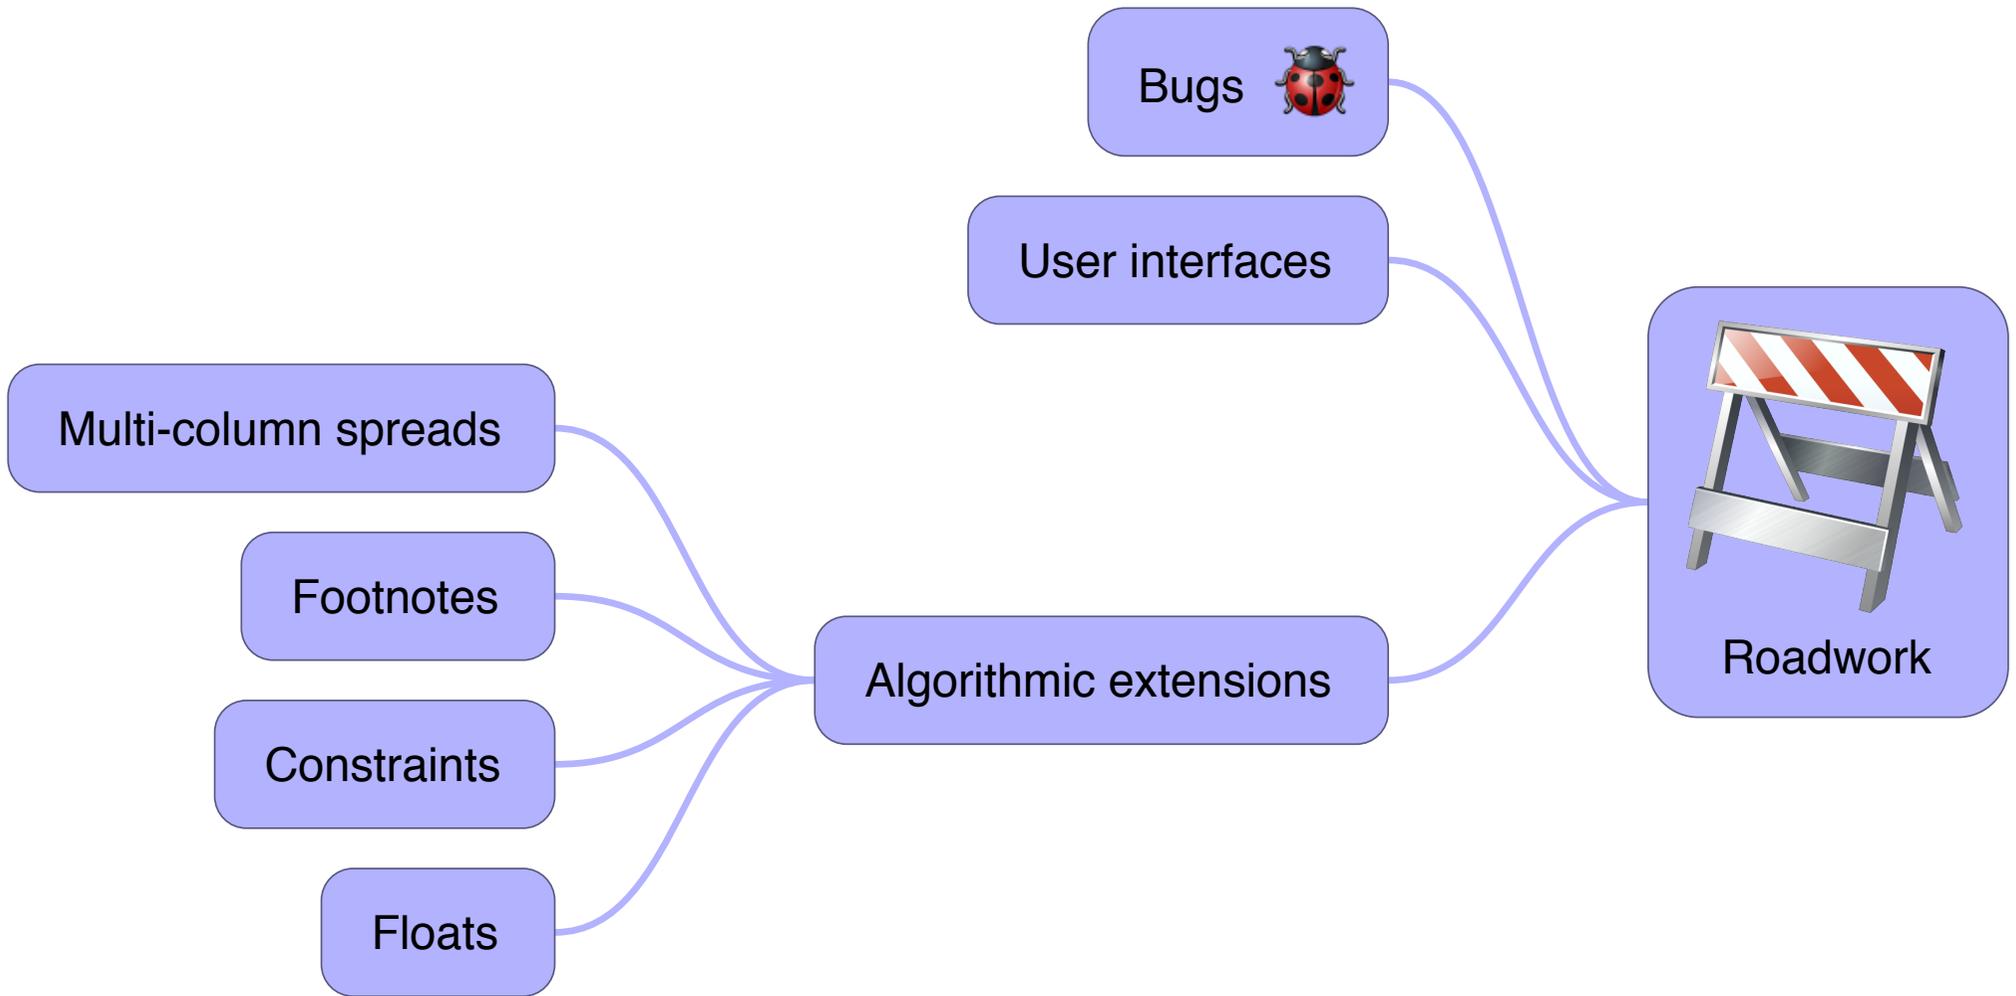
Bugs 

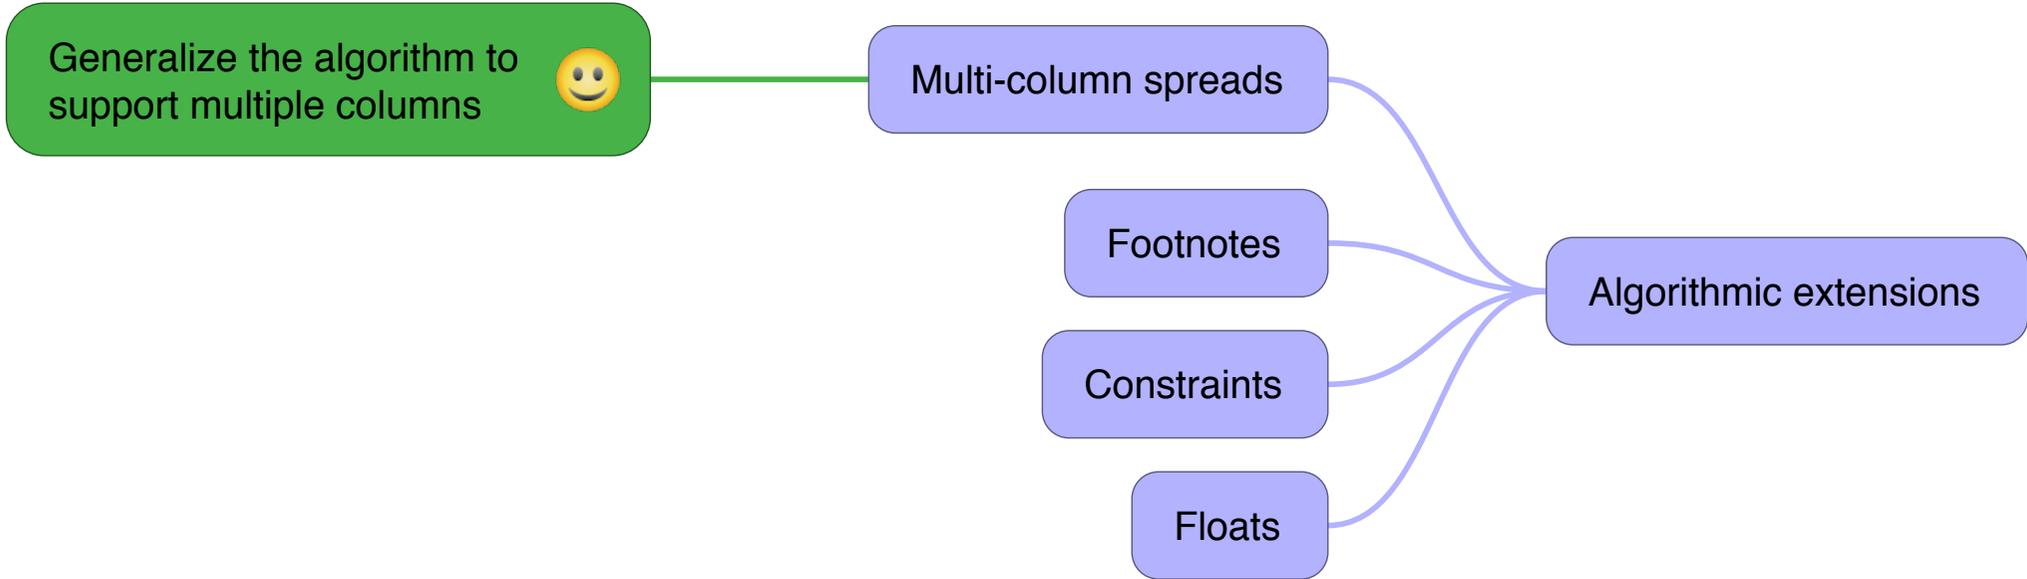
User interfaces

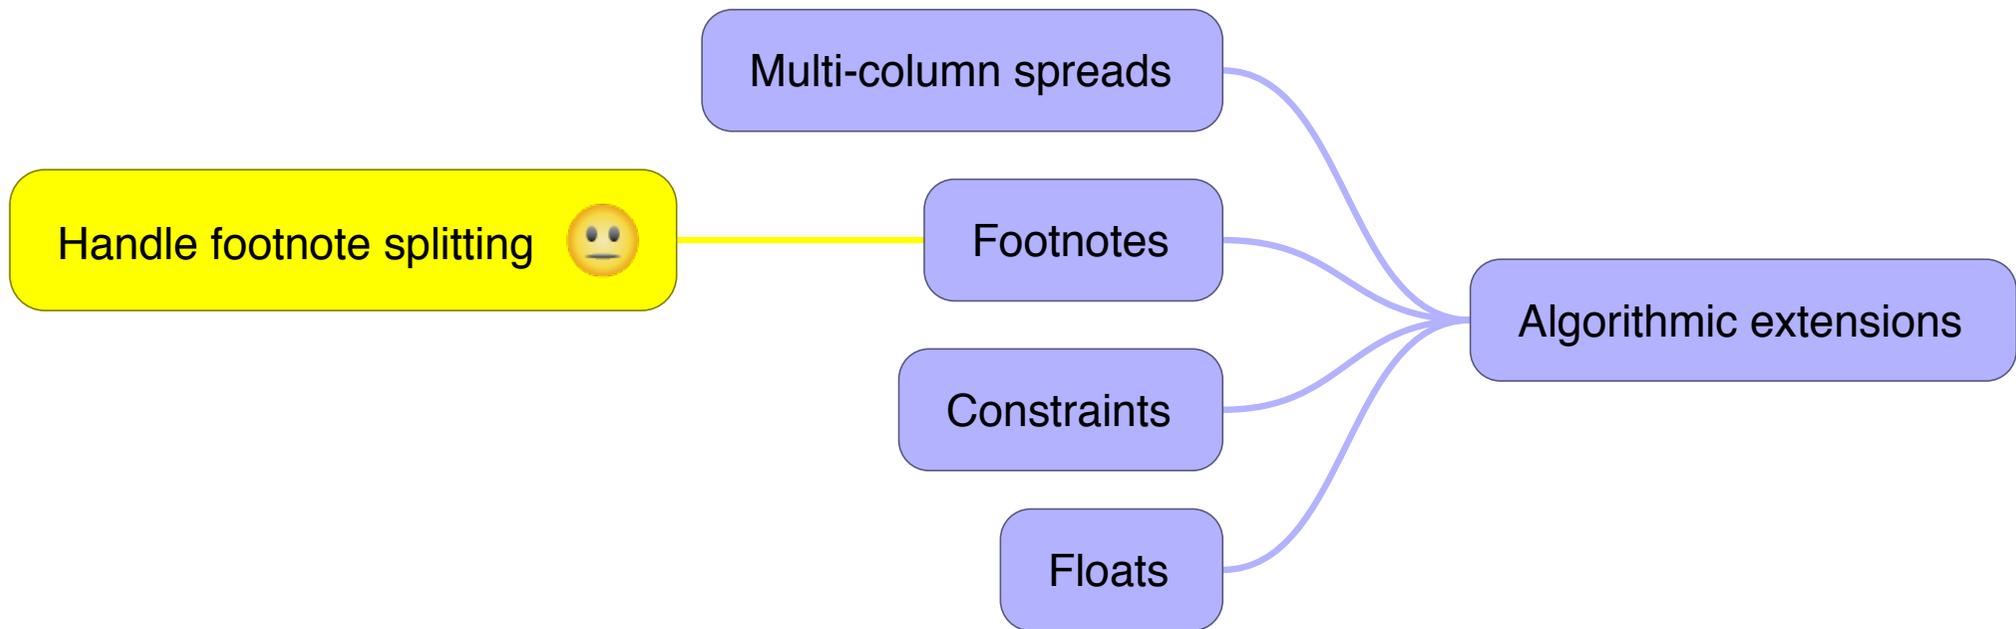
Algorithmic extensions

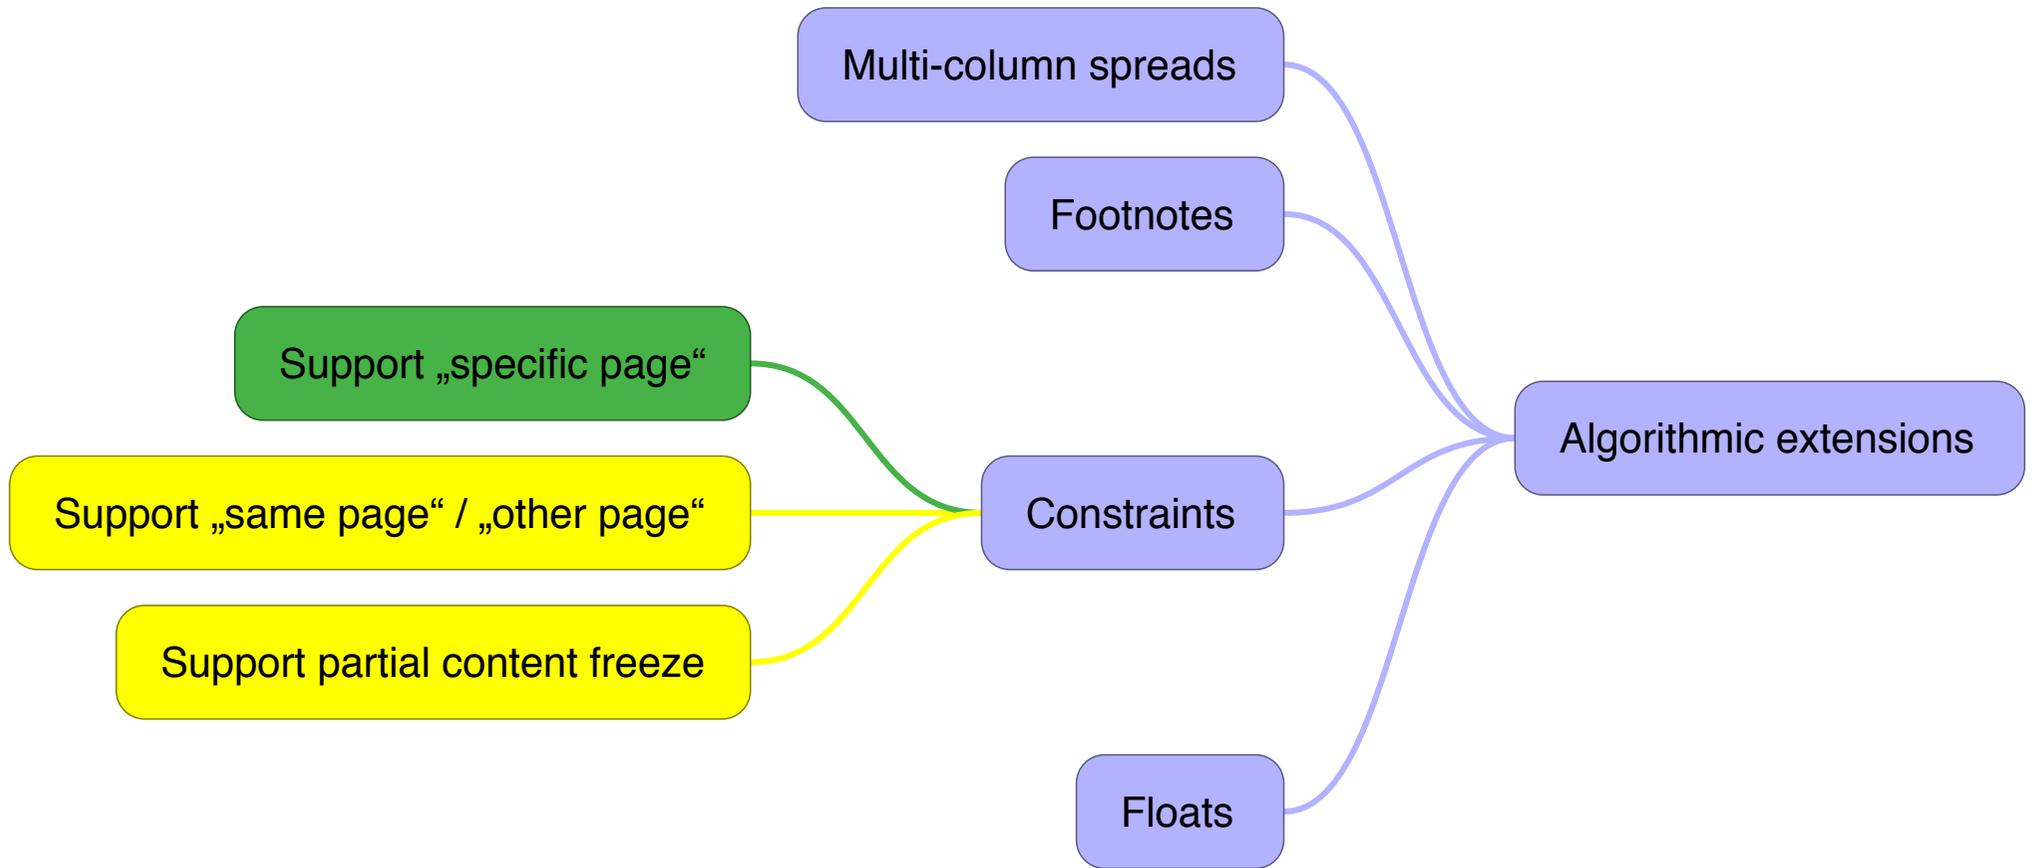


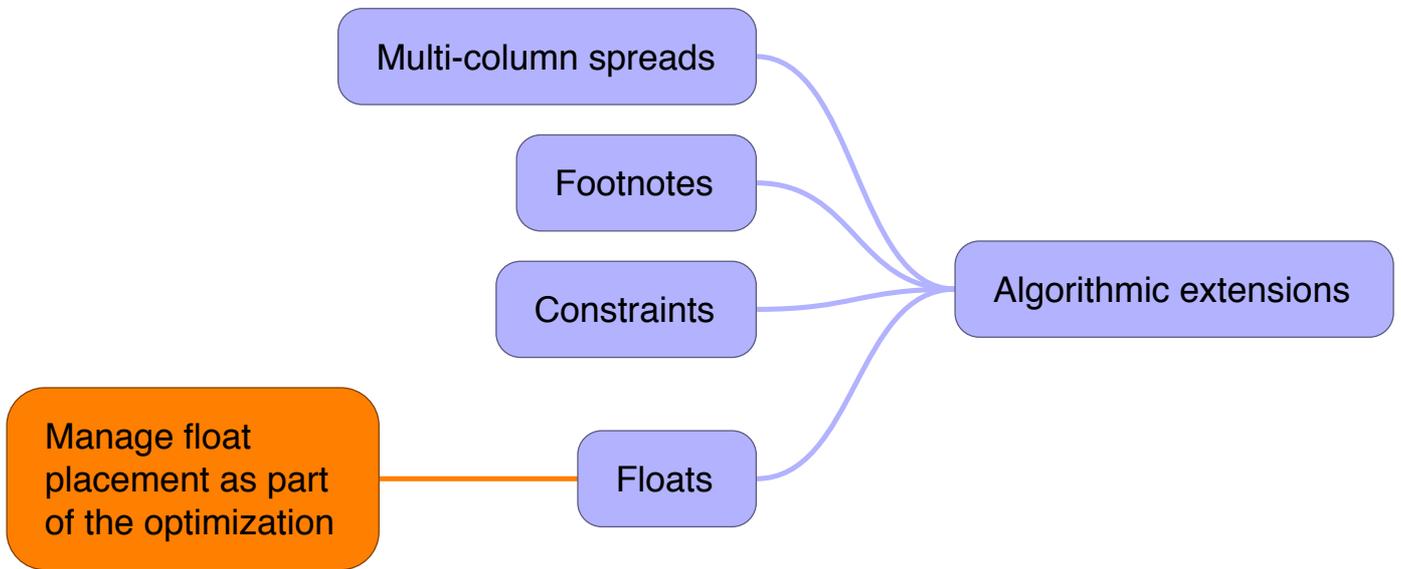












Manage float placement as part of the optimization

The number of floats that can appear on a page is bounded by a constant

Therefore number of different layouts to examine is bounded

In the algorithm with just spreads we were able to drop all inferior solutions reaching a breakpoint  $b$  when the solution involves the same page height and the same page type (verso/recto)

With floats the equivalent class is more complex: additionally we need to have the same floats being typeset on the pages up to this point

From the complexity perspective we are therefore in the same league and the time penalty should increase by a more or less constant factor compared to the algorithm presented here

thus comparable to the spread situation only that there are more possibilities

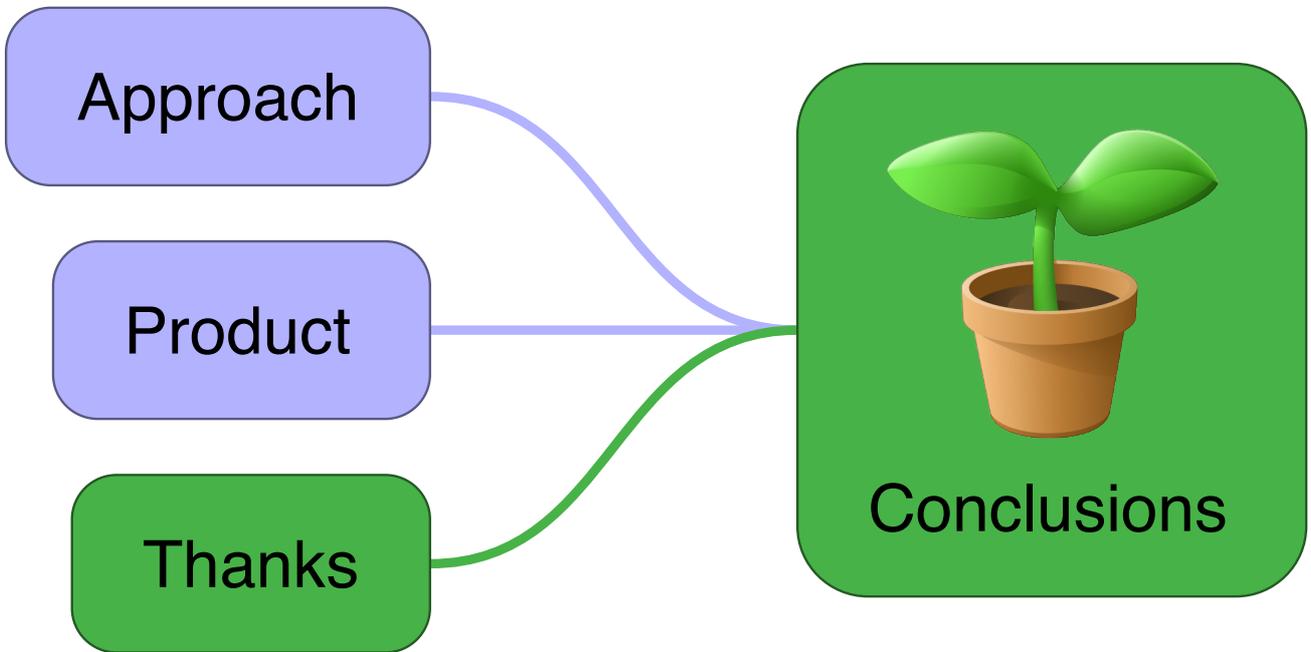
it is not necessary that they are typeset on the same pages though provided we use a sensible badness measure

Big question though: is this factor too large to make this a usable approach given current processing powers?





**Conclusions**



Depending on the job, the time penalty looks attractive

Looks usable (for certain types of jobs)



Big open question: is integrating floats possible in practice?



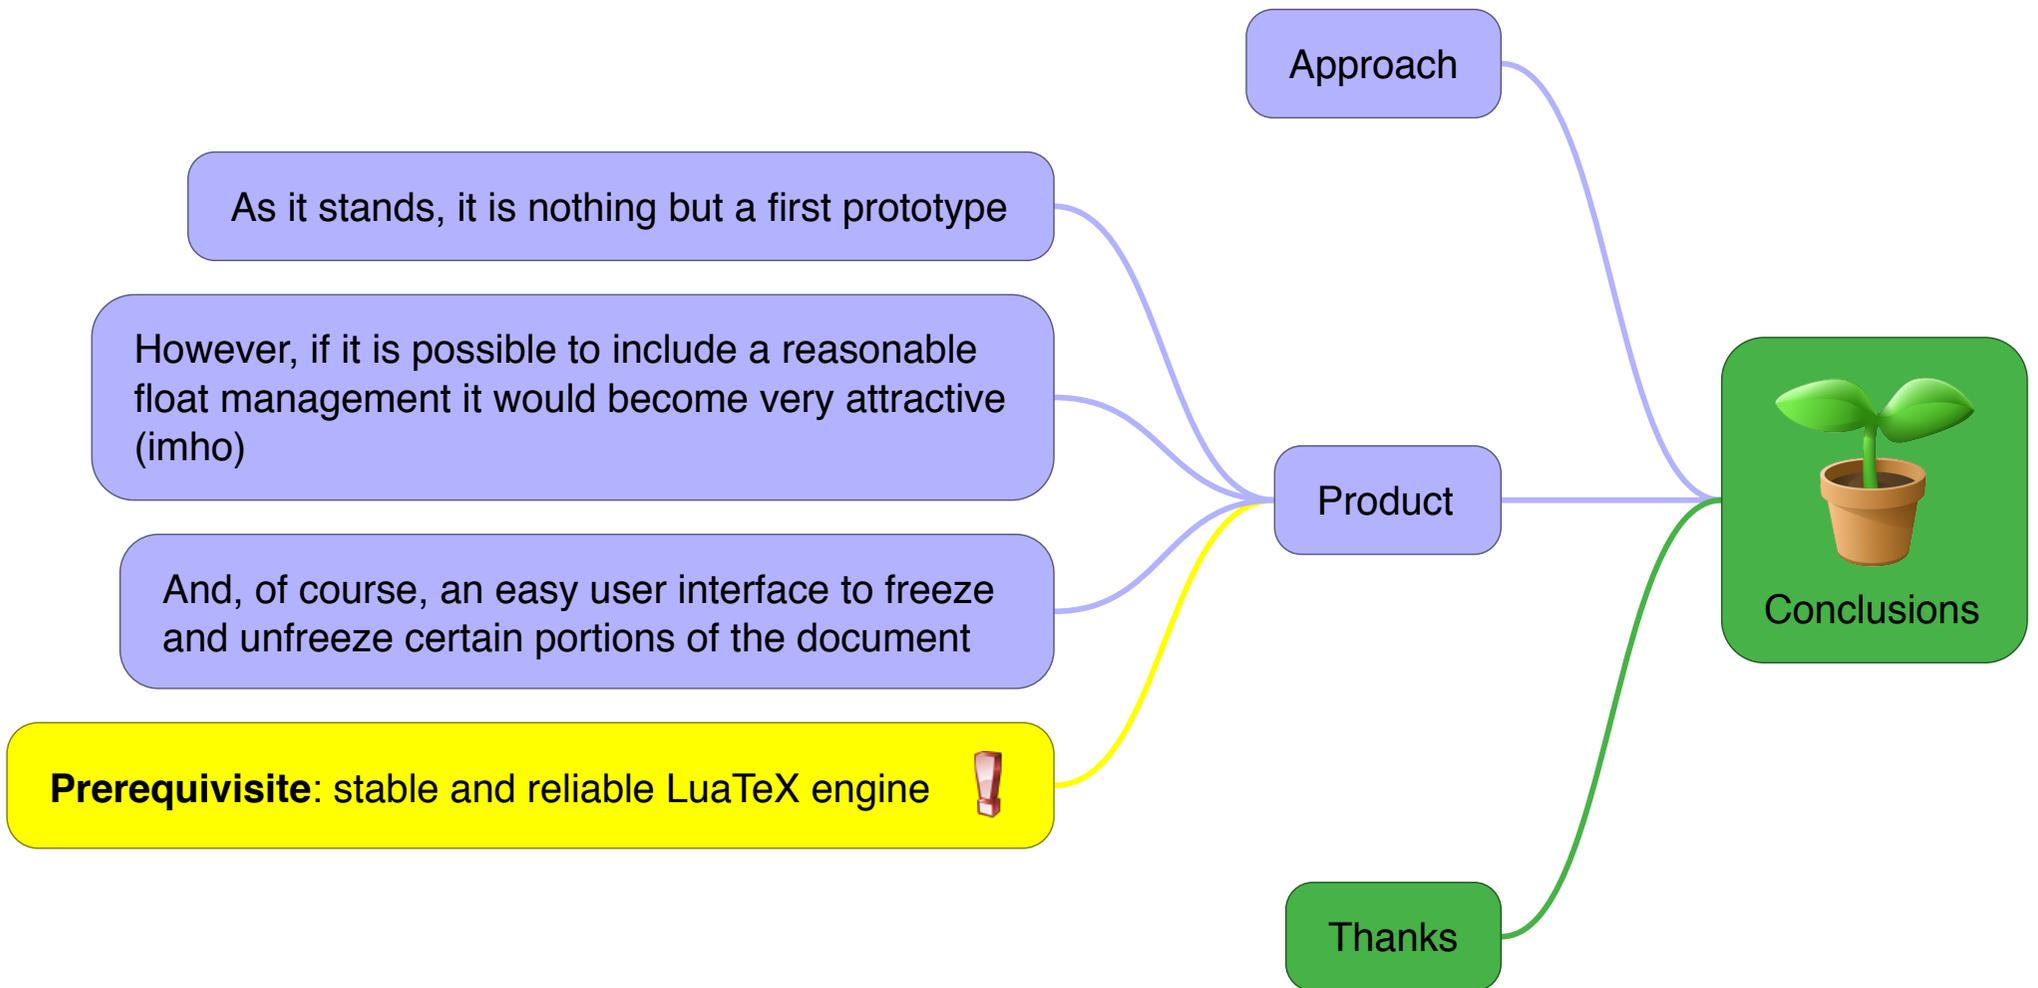
Approach

Product

Thanks

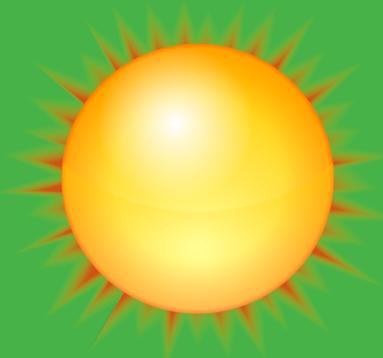


Conclusions





To **Hans** and the **LuaTeX** team for the great idea and the work carried out to marry Lua and TeX and allow me to pursue an old dream of mine — experimenting with globally optimized pagination



... and thanks to **you** for patiently listening

Approach

Product

Thanks



Conclusions

*/Alice in Bachotek/Introduction/Example: Typesetting Alice/General settings*

```
\textheight = 550.0pt (46 lines a 12pt)
\textwidth  = 229.5pt (approx 50-55 characters per line)

\clubpenalty      = 10000      % no orphans
\widowpenalty     = 10000      % no widows
\parskip          = 0pt        % no paragraph separation
flexibility
\@beginparpenalty = 9999      % strongly discourage breaks in
front of
                                     % „verse“ and similar environments
\@secpenalty      = -9000      % strongly encourage section breaks

\tolerance        = 4000      % allow fairly loose paragraphs
```

*/Alice in Bachotek/Introduction/Example: Typesetting Alice/Typeset result/Medium disaster*

- 36 pages (72 columns)
- 27 columns underfull with badness 10000 and 4 with badness > 3567 = **37% bad columns!**
- one heading (out of 11) at column start
- time to process < 1second ... time to fix **???**

*/Alice in Bachotek/The pagination framework/phase1*

The document, which consists of standard T<sub>E</sub>X files, is processed by a T<sub>E</sub>X engine without any modification until all implicit content (e.g., table of content, bibliography, etc.) is generated and all cross-references are resolved.

The engine is modified to interact with T<sub>E</sub>X's way of filling the main vertical list (from which, in an asynchronous way, T<sub>E</sub>X later cuts column material for pagination).

In particular, whenever T<sub>E</sub>X is ready to move new vertical material to the main vertical list this material is intercepted and analyzed. Information about each block (vertical size, depth, stretchability if any and penalty of a breakpoint) is then gathered and written out to an external file. If possible, data is accumulated, e.g., several objects in a row without any possibility for breaking them up are written out as a single data point to reduce later processing.

The modification is also able to interpret special flags (implemented as new types of "whatsit nodes" in T<sub>E</sub>X engine lingo) that can signal the start/end or switch of an explicit variation in the input source. This information is then used to structure the corresponding data in the output file for later processing.

The second modification to the engine is to intercept the generation of paragraphs targeted for the main galley prior to T<sub>E</sub>X applying line breaking:

- For each horizontal list that is passed to the line-breaking algorithm the framework algorithm then determines the number of acceptable variations in "looseness" within the specified parameter settings.
- For each possible variation it then does a paragraph breaking trial to determine the exact sequence of lines, vertical spaces and associated penalties under a specific "looseness" value.
- The results of each trial is externally recorded together with the associated "looseness" value of the variation.
- Finally, instead of adding a vertical list representing the paragraph to the main vertical list, a single special node is passed so that the paragraph material is not collected again by the first modification described above.

As the result of this phase the external file will hold an abstraction of the document galley material including marked up variations for each paragraph.

The result of phase 2 is used as input to a global optimizing algorithm modeled after the Knuth/Plass algorithm for line breaking that uses dynamic programming to determine an optimal sequence of page/ column breaks throughout the whole document. Compared to the line-breaking algorithm this page-breaking algorithm provides the following additional features:

- Support for variations within the input: This is used to automatically manage variant break sequences resulting from different paragraph breakings calculated in phase 2, but could also be used to support, for example, variations of figures in different size or similar applications.
- Support for shortening or lengthening the vertical size of double spreads to enable better columns/ page breaks across the whole document.
- Global optimization is guided by parameters that allow a document designer to balance the importance of individual aspects (e.g., avoiding widows against changing the page length or using sub-optimal paragraphs) against each other.

The result of this phase will be a sequence of optimal page break positions within the input together with length information for all pages/columns for which it applies. Also recorded is which of the variants have been chosen when selecting the optimal sequence.

*/Alice in Bachotek/The pagination framework/phase3/Pagination algorithm written i.../Inputs/Parameters to influence the be...*

*/Alice in Bachotek/First results/Typesetting Alice again/Adding \parskip and \baselines...*

```
\baselineskip = 12pt plus 0.05pt  
\parskip      = 2pt plus 1pt minus 1pt
```

All other values the same.

*/Alice in Bachotek/First results/Typesetting Alice again/A first demo/Setting up the parameters in p...*

*/Alice in Bachotek/First results/Typesetting Alice again/Adding \parskip and \baselines.../Results/with standard TeX*

Page 1 – second column badly spaced out (because of orphan/widow settings)

Page 2 – second column very badly spaced out (because of \* block following)

Page 7 – second column very badly spaced out (because of mouse tail following)

Page 8 – first column spaced out (because of orphan/widow settings)

Page 15 – first column spaced out (because of heading following)

In total:

- 37 pages (74 columns)
- 5 columns underfull with badness above 2100 and twice 10000 = 5,4% bad columns!
- one heading (out of 11) at column start
- time to process < 1second ... time to fix ???

*/Alice in Bachotek/First results/Typesetting Alice again/Adding \parskip and \baselines.../Results/when optimizing*

No bad pages!

In total:

- 37 pages (74 columns)
- all columns with badness < 2000 (no bad columns)
- 3 headings (out of 11) at column start

*/Alice in Bachotek/First results/Typesetting Grimm Fairy Tales/Adding \parskip and \baselines...*

```
\baselineskip = 12pt plus 0.05pt
\parskip      = 2pt plus 1pt minus 1pt
```

All other values the same.

*/Alice in Bachotek/First results/Typesetting Grimm Fairy Tales/Adding more \parskip and \base...*

```
\baselineskip = 12pt plus 0.2pt  
\parskip      = 2pt plus 2pt minus 1pt
```

All other values the same.

*/Alice in Bachotek/First results/Typesetting Grimm Fairy Tales/Conclusions/However, because of the very l...*

```
\baselineskip = 12pt plus 0.2pt  
\parskip      = 2pt plus 2pt minus 1pt
```

*/Alice in Bachotek/First results/Typesetting Pride and Prejudic.../Adding \parskip and \baselines...*

```
\baselineskip = 12pt plus 0.05pt  
\parskip      = 2pt plus 1pt minus 1pt
```

All other values the same.

*/Alice in Bachotek/First results/Typesetting Pride and Prejudic.../Adding more \parskip and \base...*

```
\baselineskip = 12pt plus 0.2pt  
\parskip      = 2pt plus 2pt minus 1pt
```

All other values the same.