

**This document contains a list of a few amazing reading and writing capabilities academics should expect to have at this time in history—we are in the 21st century after all.** Academics have access to truly portable computers capable of trillions of operations a second with high-speed Internet access and high resolution screens. However, instead of publishing documents with deep metadata and rich interaction potential, we are still mimicking wood-pulp paper rather than exploiting the potential of digital-paper. We should strive for more, as this document aims to illustrate + *show a realistic approach as to how we can achieve radical augmentations of text.*

*What is discussed here should be seen in the context of a modern laptop, maybe 13” or 16” in screen size,  
not as a science-fiction wall of text or even a 27” desktop,  
simply because this is the size of the screen of most student’s actual workstations  
and this correlates well within freedom of movement of eyes in the eye socket,  
giving a comfortable view of their work, without the need for constant large neck-muscle  
movements.*

# Document Constituent Parts : The Magic

The constituent parts of any serious academic document, or ‘paper’ have not been written or transcribed linearly in a single expulsion of speech—they have been authored—written in pieces and worked and re-worked into a coherent whole. There is a life of connections which are cut away and a world of context which is lost when the final paper is published, or in the language of most computer software; ‘Exported’ from the rich authoring environment to a generic, common format. This article highlights how great it would be to have a document of intelligent components on export, which the academic reader, be it a student, examiner or reader of a journal, could interact with to get a whole new grasp of the text. No longer just rectangular pages of black shapes of text in columns on white background, instead; text retaining their symbolic meaning. This could result in documents of magic characteristics and capabilities:

*Nothing here is impossible*

## Finding documents

Finding documents should be possible from any place something sparks the readers interest, such as coming across an author’s name and being able to instantly look up that author’s works in academic journals or in books, searching reference works for new terms or translating foreign text.

Collections of documents should avail themselves to easily navigating the documents and changing the views to allow for further discovery of materials:

## Collections of documents

Collections of documents could be analysed and manipulated as true hypertext nodes, arranged based on their authorship, dates, keywords, what names were mentioned, who they cited and in what order, and who had cited them, or any other ways the user prefers.

External evaluations of the document, through machine learning, citations or other, could be presented in a way most useful for the reader, such as through a whimsical shuddering animation of suspect documents or shining highlights for documents deemed by the system to have been cited the most, and most favourably, which the reader has not read

yet. Perhaps. Or perhaps they will simply be shown in lists—whatever truly suits the needs of the multiple-document reader at the moment.

The collection of documents could be colour coded based on contents in a way the reader is familiar, with connections shown based on pre-set preferences and instantly changed on a whim to give further insight. Perhaps documents authored by someone the reader has not read yet will be red. Maybe long articles have one colour and short articles (in the conference sense) another colour. Maybe it could be presented as a magical Harry Potter library or that of a futuristic sci-fi display. Whatever design, use and re-design will show to be most useful and pleasant to use.

### Pull forth any individual document

The reader could pull forth any individual document from such a collection, and interact with it, getting more information while keeping the rest in context:

The document pulled forth could appear in any way the reader prefers, including showing only the abstract, with the last sentence highlighted—often a useful indication of what the document really is about, as Chris Gutteridge showed me. The reader could ‘unfurl’ the document to see headings and keywords and then choose to see citations in context to get a picture of what the document refers to. Or not, simply seeing any pictures in the document or accessing any videos—or whatever the reader wants.

Ted Nelson style lines could connect citations to sources visually.

If interesting, the reader could easily tag the document or snap the document into a document view where all others fade away for now, giving the reader an optimised reading experience:

### Single document view : basic options

In a single document view the reader can study the document in the traditional, analog way. If the reader prefers a traditional view, this should not stop the view from being formatted to whatever academic style is desired, such as numbers for citations (brackets, hard brackets and superscript all supported) or as author-date, no matter what the author specified. References should be in alphabetical order or as they appear in the document (with or without headings from the document also shown in the References section for context). The text should be encoded in such a way as to entertain any view.

## Single document view : advanced options

If the reader prefers to study a single document in ways only afforded by using a powerful, 21st Century computer, the only limits should be the intensity of their curiosity and the depth of their literacy. A deeply literate digital text reader should be provided with the tools to truly get to grips with the text, no less than a fighter pilot, athlete, racing car driver or gamer has advanced tools to perform their tasks. I would argue that tools for thoughts are at least as important as any other tools we can develop. The content of the document should most certainly not be the constraint to the development of powerful interaction tools.

On opening, a single-document ‘view preset’ can be set to open documents in, such as highlighting terms which are important to the reader. This is much like a photographer can specify what treatment photographs should be given on import in an image editing program. This view preset should be possible to change at any time, and the user should be able to change the view at any time—it’s only a convenient start.

The document could be folded and unfolded to give different views of sections based on headings, keywords or other.

Salient details, such as the shape of citations—where they appear in the document, if they are repeated, and if they point to known or unknown documents, should be easy to see.

Instant access to sources could provide deep insights with minimal mental load. There is little reason why they should be buried in the Reference section and not digitally connected to their appearance in the text—and not just as a one way link—but with a click to spawn an informative and interactive pop-up displaying all the relevant citation information.

Any word or phrase could be instantly searched for or looked up in any database, with the smallest mental effort, as with all the interactions discussed here, performed as smoothly as a world-class downhill skier navigating an advanced slope, with the greatest joy and ease\*.

Seeing all the occurrences of certain text should be a quick-click, and just as easy to dismiss, not just having to go through a document hunting for yellow-highlights as is currently the fashion for Find.

Any text further elaborated by the author as a term in a glossary can appear without the reader even explicitly asking for it, by simply asking to see all the occurrences of the text in the document and any associated glossary reference is conveniently shown prominently at the top of the results.

Graphs could be spun based on any criteria and manually or programatically laid out to show any number of relationships.

Navigating a document through scrolling could provide a completely different view

from reading, by for example having body text shrink to give more prominence to headings and pictures taking the place of names. Icons could replace company names, flags could replace country names and portraits replace names and references to people. It could be useful for this to automatically happen when scrolling through a document, for quick visual recognition, then fade back to text when the navigation is done. Reader designed colour-glossaries could be used to colour keywords in the document on scrolling, such as yellow for text about people, blue for concepts, green for technologies and so on, making it clear for the reader what sections cover what topics. Doug Engelbart and I experimented with this based on his suggestion, finding it surprisingly useful, but it needs to fade when navigation is done—unless the reader prefer it stay.

Documents specifically prepared for students could have curated views where the student could be shown certain highlights and certain interactivities could be included, such as multiple choice questions on pages which only appear after a certain time period so the student might as well read the full page first.

## Annotating

Annotations should be as involved as desired or as simple as preferred, yet the context should remain rich, with colour coding, grading, drawing, whatever the reader prefers to do.

For example: Select text and ‘r’ for red will mark the text read, ‘g’ for green and then easily view all the marked text of a specific colour in one document or a collection.

After annotating, the reader should be able to easily find text based on what was highlighted (searching for only highlighted text is currently not possible, which is surprising) or what was written by the user on a page.

Not only that, all information about the annotating should be retained, including where and when it was done, to allow the user to later search based on time and location, such as ‘show the text I annotated when at home last week’.

## Citing

Forgive the play on words but citing could be made to be more exciting.

Robust citing should be a matter of simply copying and pasting and what is pasted is a full citation. This is one of the functions we have been able to implement. More on that later.

Pasting a DOI or a BibTeX reference into a document should be resolved into a full citation as well.

Citations should include the full cited document when there is a concern about the document disappearing and where suitable rights are available.

It is not unreasonable for the cited text to include all the references to the cited document to help the reader see how it is situated in the information landscape.

Link types should be possible, to make it clear if the citation is supportive, informative or in disagreement.

## Authorship

When authoring a new document it should be possible to instantly toggle between different views of the document, including ‘pinching’ to fold the document into an outline. When folded, any highlighted text should appear in-situ under headings to remind the author what sections need work.

An easy keyboard toggle should turn on a focus mode where all text except for the paragraph the user is working on fades. The point is, the document should be manipulated like a trained sculptor manipulates clay.

In terms of views, any names could show up on command or any bold text could appear. Whatever is useful.

Sections should be easily marked as ‘done’ so that in a long document the author doesn’t need to be distracted by what is not being worked on. Other ways to fold and manipulate the view could be to hide everything in a section that is not bold, or collapse lists. Of course, in a default view this would all be visible, as well as on export, unless the author chooses to keep them hidden.

How about writing ‘auditions’? This is a term and idea borrowed from Apple’s Final Cut video editing software where the user can make an edit in a section of the timeline and then save that edit and make a new edit for that section, and later choose which is preferred.

And how about being able to cut any text and know it will be kept in a a Cuttings list, so you can always paste what you might have cut and forgotten to paste?

A non-linear thinking space of a manually and computationally laid out graph should be instantly accessible and connected to the main ‘word processor’ view of the text. This view could accommodate keywords, headings, citations and more, in unison or separately. Timelines should be possible to apply, as well as connections and layouts based on any meta for any of the nodes.

Margins could be used to magically provide note spaces and links to the document to extend the author’s memory and perspective, as mini-versions of a non-linear thinking space.

It should be possible to write computational text as easily as prose, including adding live snippets of code or simply text that can change depending on context.

Computational text could enable sentences to change based on context, to always be up to date for example. It could also allow the text to be interactive, to allow an author to write an equation or a temperature and the reader can click to change a variable or convert the temperature to another format.

Text Expanders used to be a big thing, they seem to be less common now, but can be powerful: Easily add links and descriptions based on a few words. This is a bit like the suggestions we have available for typing, on handheld devices and computers, but with pre-programmed options as well as some intelligence, such as the current ability to automatically add an address, but expanded to provide live suggestions such as ‘King George I was born on’ and the system giving a suggestion for the next text to be ‘1660’. Every sentence typed could then be treated as a search command (to Google or Wikipedia fx.) and action command (to Siri fx.) so that typing something as simple as ‘as I wrote last week’ could allow the user to have a search of all documents typed last week. It could also be implemented in a way where on command, the system would take any text to the left of the cursor and send it to our Liquid\* interface for instant actions including search. It’s not that selecting text (which is the normal use of Liquid) takes time, it’s just that allowing for smoother actions means it’s more likely the user will use them.

Active Templates could help the students write a success full assignment or any writer an article or book.

## Organising your own documents

Organising ones own writing is currently a challenge for many, including remembering what one has written before and how to link to it to reduce repetition. With smarter documents and tools, imagine being able to search your own published documents to refer to them—or your notes—at the time of writing or at the time of thinking, by keywords, when you wrote something, or any other criteria, in order to better spin your own work into your web of publications. This could take the form of simply having a keyboard shortcut handy to search your own blog or the journal of your field so that you have instant access to a search view to further refine your search and insert a citation.

Keeping the Open Dialog and the Export Dialog separate would save many a click. Usually we don’t export to the same directory as our source documents are, but in most software we need to constantly change directory when opening and exporting. A detail to be

sure, but a lot of clicks and mental time can be saved by changing this. Adding a more useful listing view in the Open dialog, including being able to search by tag and directory (not one or the other as in the current version of macOS) would further allow for basic organisation.

Another wrinkle could be to let the user specify what kind of document their working/source document is, when saving the document, such as final thesis, notes, drafts and so on, which would be useful in dynamics views and searches (such as above) but which could also provide a visual ‘frame’ for the document in the Open dialog and on the desktop, to allow for quick glanceability: That’s my big document, those are scraps, those are published papers and those are contributing research notes...

## Desktop

This is out of scope for developers of applications but it’s worth mentioning since the desktop is ubiquitous in our work. The operating systems’s natural, shared space, has great potential to be more useful, such as subtly showing the size of documents by mimicking a thicker stack of ‘paper’. Attached notes can appear as little stickies attached to a corner. Older documents can yellow slightly like fax paper or parchment. And more, much more.

Further questions include why can the desktop on our computers not have tabs? Imagine layouts for different projects with different folders open and dynamic views showing external or internal records by whatever criteria we choose? On a macOS MacBook these tabs could even be accessible through the Touch Bar.

Why not have a built-in graph view of all our documents? 3D views have been experimented with in the past. With the power of a modern laptop we could provide incredibly powerful views for the user to navigate their own documents, how they relate to each other and how they relate to the rest of the world of documents.

## Calendars

You should be able to use your calendar and request to see pictures you took or documents you wrote during: Just click on an event or drag-across a time-span and a beautiful array of anything associated appears, including transcribed video and audio records of conversations, as well as emails and other messages, but if that is too much, click to specify only those from the morning, or only those from sunny days; after all, you might have a vague recollection of the circumstances around a conversation or a note and time is a powerful dimension in navigating knowledge. The notion is of a time browser where your system surveils all you do,

giving you easy access to the past (as well as planned events in the future), to when documents, web pages, news, social media posts were read. To when documents were published and so on.

There is no reason you should not be able to perform, visually, via a spoken assistant or whatever other means, a query to instantly have a list of all the times a keyword was spoken in a video conference over the last week, with certain people, and when you get the result you can easily expand to see further parts of the conversation. It should also be possible to cite any of the parts of the conversation, provided they were marked by the participants as giving permission, into documents where the reader can reconstitute the whole meeting should they have access to the recording.

## Locations

Locations should be tracked so that you can easily search for that webpage you read in the evening last week when in town, for example. You should also be able to track what wrote at different locations in a document, where the document tags every sentence with location and time data.

Referring to locations in writing should allow for tagging of different ways of identifying the location, from full place and country name to GPS and what3words.

## People

Much of what we care about and write about is people. We have lists of people and information about them in our address books or other systems. Apple has had what they call Data Detectors and this includes people-detection to some extent, so how about being able to click on a name anywhere and see all correspondence with that person, anything they have published, their personal website, links to what they have said in conversation with you through online conferences and more?

Historical or public figures should have biographical information attached for alternative document views in timelines, geographical views or relationships views, or other, using Wikipedia/Wikidata or other data, including that from a curated glossary.

This comes to the heart of what Doug Engelbart called ‘symbol manipulation’: Text conveys intent and when referring to something as important as a human being, there is a lot of information that’s available and pertinent and this should be made as available to author and reader as richly as possible. Symbols mean something, we should allow the users to

interact with the meaning, the intent of the author, as richly as possible, by allowing the connections inherent in the symbols to be as interactable as possible.

## Voice

Voice should not be excluded but extended into as powerfully as possible. Though voice is no substitute for a keyboard or cursor when working, should be deployed to give you opportunities to speak to your watch (for example) with commands such as “add a note to my most recent document saying ‘I need to add more historical background’” and it should appear on that document as a yellow sticky note, in a top corner.

When reading, speech to text should work as conveniently as possible, with spacebar to speak and spacebar to pause (currently spacebar starts back at the beginning rather than resuming when tapped again), as well as providing an option to record the speech to a podcast which the user can listen to, with markers so that the user can see how much then listened to when they return to the document, something like Amazon’s whisper sync.

When reading, the user should also be able to add voice notes, either to the open page or to a specific section. This should be transcribed as well as left as voice, for easy teacher comment on a student paper, particularly or for a student to comment on the research they come across.

Voice was once far beyond the reach for a small developer such as us, but with powerful ML processing capabilities and Libraries, it is time to extend into other media to connect it usefully to text.

## Images

Images should certainly not be ignored. In the way that HTML image maps allowed a designer to specify what interacting with specific parts of an image allowed, ML should allow parts of images to be tagged (and manually adjusted by a producer should there be a need), allowing the user to interact with a series of images purely by clicking on sections. As I suggested to a professor in Germany who was interested in downed WWII aircraft pictures, why not make it possible to click on the aircraft and get as much information about the model and specific airframe as available, and then click to see more of that type, for example, and click on the sky and ask what the weather was on the day the picture was taken or the day the aircraft was downed. And so on. This data could be stored in an extended EXIF or as information printed on the frame of the image, as notes on a contact sheet perhaps.

It should be possible to link inside a video in the text, which is still displayed inline, for quick navigation, not just when played externally.

Maps with coordinates available to the document should be embeddable so that the reader can choose where to see the location; in a system map or a specific application of their choice, or in the reading application if it supports map views.

Tufte Sparklines should be embeddable for live and historic data.

Importing of any image should OCR the image to extract any text and allow this to be used for searches and views.

## AR, VR, multiple monitors

AR, VR and multiple monitors can increase the user's available workspace to allow them to build something like the 'murder walls' seen in crime TV shows to organise data in rich fashion, but for text, it will remain crucial to have a high-resolution, great quality screen of the kind of size laptops have today, to make close reading pleasant. The work we focus on with Augmented Text is this intimate reading and writing space though connections to future spaces around the users room clearly has interesting potential.

## AI, ML

Artificial Intelligence will have profound effects when truly realised, in ways we likely cannot imagine today. Machine Learning however, can have a profound effect on text interactions today, with analysis of documents based on grammar and contexts to augment our ability to handle large corpuses as well as to better grasp the intentions of individual documents.

Constant ML analysis of the document could be used to give the user suggestions as to what is missing, depending on the type of document, what is repeated in the document and even provide contextually relevant questions for what to write next.

Simple things (for a human), such as natural language processing can produce more useful 'Find' results, if, for example, a person's name is search for and in another sentence she is simply referred to as 'she', or 'her' and the system sees that it is the same person.

Sentiment analysis can help analyse student papers to help the teacher spot students who might be having a hard time for example.

Writing with ML could give us semantic autocomplete options, where we write what we can and the system suggests what to write back, based on a 'knowledge' of our domain

and our writing style, so that we in effect have magic pens to sometimes guide, not always have to 'spell things out'. This is related to the Text Expanders function outlined above, and could produce a host of interesting implementations.

## Smaller improvements

Many smaller improvements, such as using the ESC key to go in and out of full screen and well designed keyboard shortcuts and gestures can further provide a more pleasant and deep interaction over time. Polish makes a difference.

## Integrations

It should be possible to copy text from a thinking environment like Roam, Notion or Em, into a word processor and when exported the full connections of the thoughts can be expanded. Same with graphs and other specialist applications.

## Publishing

When publishing a new document everything the author would like to have retained of connections, context and semantic meaning will be retained, yet have document which will stand on its own.

When opening a document as a teacher or examiner, the document could be processed to check for plagiarism and reading level etc. and even checked if specific canonical texts in the field have been cited, providing the reader with either a report if significant issues have been found, which the user can click a button to send to the student, or simply opening the document with perhaps a few highlights. A summary could also be generated and grammatical issues highlighted.

Students should have access to such processing on opening, including the ML summary, so that the student can see if the document expresses what the student think it should, before submitting. Such a process could be valuable when submitting to a journal as well.

## Useful Or Utopian?

Some of this, dear reader, may grab your interest and some of it may simply seem fanciful, as though just a special effect in a movie like *Minority Report*. What I hope you can agree with however, is that there is powerful potential in unleashing text interactions. A cyberspace for symbols.

You may have completely different ideas of how it should be done, but I expect you may at this point also share a frustration of why this is not currently possible. It is not simply the near monopolies of a few large companies (Apple, Google, Facebook, Adobe and Microsoft) and the limited resources for the small companies who build alternative reading and authorship tools. It is also because the medium of exchange, the very store of our knowledge is inadequate:

# Frozen Documents

The reason we cannot have such richly interactive documents is because PDF, which is the default format for academia, simply does not allow it. Even Adobe, the developers of PDF admit this, which is why they developed a ‘Liquid View’ for PDF using—get this—machine learning!, to extract some basic formatting.

Though competing document formats come and go, PDF is too entrenched to be replaced easily.

It should also be appreciated that PDF does a tremendous job of freezing documents, which is crucial for academic documents. It is important not to throw the baby out with the bath water, PDF performs one academic function very well.

A multi user, multiple version, editable document such as Google Docs would not suffice for academic analysis and citing since it is very messy to cite something which may change.

The design constraint to overcome this then becomes allowing for robust freezing of document contents *and metadata* which can allow for rich, fluid, interaction.

# Digital Fragility

Digital text introduces inherent flexibility to allow for choosing its appearance and to allow for copying, undo, spell-check and so on. However, this flexibility is also a limitation.

While digital text allows for instant connections through mechanisms like hyperlinks, these come with the inherent brittleness of lacking a permanent physical substrate. If a physical document is placed in different libraries, offices and homes, it is definitively harder to track down a copy if you don't have one, but by presenting the bibliographic information to other people you will be able to locate a copy to purchase, borrow or copy, even if it was not at the address you were first sent to.

If the document is only digital then it disappears if the server goes down or if DNS rent has not been paid and the address becomes invalid, you loose access to the document.

Furthermore, the digital text can be rendered unreadable if the software reader tools, or the platform it runs on stops being available.

# Meta-Solution

There is a way to allow frozen documents—the very currency of academia—to exhibit all the fluidity a deeply literate reader can need—*written in the text itself*: We call it Visual-Meta.

Visual-Meta is at the most elementary an Appendix at the end of the document with BibTeX formatted citation information so that the document ‘knows’ what it is: who wrote it and when, as well as what the title is. The Visual-Meta can further contain formatting information such as where the headings are and what glossary terms are used and so on.

This is only the very beginning of what is possible when the document becomes an active ‘neuron’ in the distributed academic ‘brain’, capable of acting as a self-contained document but deeply connected and deeply open. This illustrates how depth comes from surfacing, something applicable both to information and the interactions we have to manipulate the information.

In academia the visual styling of the text is important. However, the visual style of the document should not overshadow the huge benefit embedding meaning as well as focusing on the presentation can achieve.

Visual-Meta is, to a standard PDF viewer, just an extra Appendix of a page or two. It looks the same information as you will see on one of the first pages of a book (title, author, date, publisher and so on), but it has the power to transform the interactions possible for the reader as well as how the document is displayed, including such niceties as changing the academic styling to suit a different discipline, at will.

Documents today are a mess of frozen data with a severe lack of depth, and we are in the eye of a hurricane of information overload and fake news—however; **we can, quite literally—write our way out.**

## Sticking with PDF

Of course, PDF documents contain the same brittleness that all other document formats do in that they require software to be read. However, they are the de facto standard for so many documents which means that the likelihood of support being continued for a long time is high. Additionally, a benefit of Visual-Meta is that the documents can be printed out, stored, scanned and OCR'd and no metadata will have been lost.

# Building Augmented text Tools

So far we have designed the basics of the Visual-Meta system, as outlined at [visual-meta.info](http://visual-meta.info) and implemented it in the Augmented Text Tools Reader (PDF viewer) and Author (word processor) for macOS [www.augmentedtext.info](http://www.augmentedtext.info).

It is also employed in our published book *The Future of Text* [futuretextpublishing.com/future-of-text-2020-download/](http://futuretextpublishing.com/future-of-text-2020-download/)

I invite you to have a look at the links and see-and what this looks like and, if I may be so bold, download and experience how far we have gotten in making all this real, and imagine how we can progress ever further.

Not all our work is enabled by Visual-Meta, but it is what allows advanced features to be retained when documents are made public, for further advanced interactions in other software.

We are building a community of augmented text tool users as part of our annual Future of Text Symposium [futureoftext.org](http://futureoftext.org) and through the contributors to *The Future of Text* book, which includes encouraging students to think more about their text tools through the annual competition to write a contribution for the book.

Thank you for your time dear reader. I wish you may read similar words in much richer ways in the future.

Frode Hegland

Norway, December 31st 2020

*frode@liquid.info*

*\* Instant access to searches, references, translations and conversions is what our 'Liquid' tool enables, our first Augmented Text Tool, growing out of the concept of 'Hyperwords', which Douglas Engelbart so very much appreciated: <https://youtu.be/UrBi4h8jT54>*

# Visual-Meta

```
@(visual-meta-start)

@visual-meta {
  version = {1.0},
  generator = {Author 6.1 (907)},
  }

@article {
  author = {Frans Alexander Hegland},
  title = {Weaving},
  month = dec,
  year = {2020},
  institution = {University of Southampton},
  }

@headings {
<json>
[ {name="Document Constituent Parts: The Magic", "level":"level1"},
  {name="Finding documents", "level":"level2"},
  {name="Collections of documents", "level":"level2"},
  {name="Pull forth any individual document", "level":"level2"},
  {name="Single document view - basic options", "level":"level2"},
  {name="Single document view - advanced options", "level":"level2"},
  {name="Annotating", "level":"level2"},
  {name="Citing", "level":"level2"},
  {name="Authorship", "level":"level2"},
  {name="Organising your own documents", "level":"level2"},
  {name="Desktop", "level":"level2"},
  {name="Calendars", "level":"level2"},
  {name="Locations", "level":"level2"},
  {name="People", "level":"level2"},
  {name="Voice", "level":"level2"},
  {name="Images", "level":"level2"},
  {name="AR, VR, multiple monitors", "level":"level2"},
  {name="AI, ML", "level":"level2"},
  {name="Smaller improvements", "level":"level2"},
  {name="Integrations", "level":"level2"},
  {name="Publishing", "level":"level2"},
  {name="Useful Or Useful?", "level":"level1"},
  {name="Frozen Documents", "level":"level1"},
  {name="Digital Fragility", "level":"level1"},
  {name="Meta-Solutions", "level":"level1"},
  {name="Sticking with PDF", "level":"level1"},
  {name="Building Augmented text Tools", "level":"level1"},
  {name="Visual-Meta", "level":"level1"}
]
</json>
}

@(visual-meta-end)
```