Improving Human Comprehension With Individualised AI-Based Text Modification

Nikolas Klein
Hochschule für Gestaltung
Schwäbisch Gmünd
hello@nikolasklein.de

Florian Geiselhart
Hochschule für Gestaltung
Schwäbisch Gmünd
florian.geiselhart@hfg-gmuend.de

Christoph Labacher
Hochschule für Gestaltung
Schwäbisch Gmünd
mail@christophlabacher.com

David Oswald
Hochschule für Gestaltung
Schwäbisch Gmünd
david.oswald@hfg-gmuend.de

Florian Ludwig
Hochschule für Gestaltung
Schwäbisch Gmünd
info@florian-ludwig.de

Abstract

In order to improve human comprehension of information the problems of “Information Overflow” and “Unstructured Information” have to be solved. In this paper suggestions how to solve these problems are made and the capability of Artificial Intelligence and Individual Media to improve comprehension is examined. Two conceptual applications demonstrate possible use cases and describe how future technology to summarise and augment information can be applied in a human-centered way. The ramifications of these technologies are discussed with the intention to spark a discussion about the impact of AI on human information comprehension.

Introduction

Computer scientists and designers alike try to build tools to help people solve complex problems. In order to solve problems, gaining comprehension of the available information is crucial. As Vannevar Bush noted: “A record if it is to be useful to science, must be continuously extended, it must be stored, and above all it must be consulted.” [2] This paper presents two new approaches how existing records can be “consulted” in an improved way.

In 1962 Douglas Engelbart proposed his framework for “Augmenting Human Intellect” [4]. He mentioned three capabilities of humans that needed to be increased: “to approach a complex problem situation, to gain comprehension to suit his particular needs, and to derive solutions to problems.” In order to improve the way we gain a suitable comprehension, according to Engelbart we need: “more-rapid comprehension, better comprehension, the possibility of gaining a useful degree of comprehension in a situation that previously was too complex”. Today these challenges still exist: Traditional Knowledge Management (KM) is struggling to fulfil Engelbart’s requirements: According to Fischer and Otswald one of the main problems of modern KM is still “Information Overload” [5] – standing in the way of more rapid comprehension. Diao et al. additionally mention “Unstructured Information” as a problem [3] – standing in the way of better comprehension by not putting the necessary information within the user’s reach.

There have been propositions made on how to use Artificial Intelligence (AI) in order to improve the way knowledge is turned into digital information, classified and searched [3]. Birzniece surveys multiple ways in which AI has been used to address challenges in KM, such as using expert systems, data mining or case-based reasoning [1]. However those experiments focus on improving the technical infrastructure of KM. None of them use AI to improve the users’ comprehension of information. Our approach is to look at recent breakthroughs in Artificial Intelligence, like the Abstractive Sentence Summarisation method shown in [14] and Natural Language Generation, and consider how these might help achieve improvements in the comprehension of information and advance knowledge acquisition technology in the future when developed further. With a human-centered design approach we put our focus on how these technologies need to be applied to be accessible to users and find acceptance.

In the remainder of this paper first, we introduce the notion of a new kind of media, called “Individual Media”. We present two conceptual applications, created using human-centered design methods, which attempt to provide rapid and better comprehension by using AI, while retaining the users’ control over what they retrieve. We discuss the impact of “Individual Media” and explain what we believe is important for engineers and designers to focus on when developing similar applications to those presented. The goal is to spark a discussion about the dangers and opportunities of such new technologies in KM and encourage the creation of guidelines for people working in this field.

Individual Media

During our research into the history of media and information consumption we identified a trend that will likely have great impact on how information is created, consumed and shared in the future.

Looking at the development of media through time (fig. 1) we see two incisions that significantly changed the character of media. Coming from spoken language and oral tradition, which required a direct interaction between “creator” and “consumer” of the information, over handwritten documents, the first incision is the invention of the printing

Figure 1: A timeline of media innovations and incisions

press. It marks the dawn of mass media. After this invention there was a one-to-many relationship were there had been a one-to-one relationship before: one author being able to communicate exactly the same message to a large, diverse audience – even throughout time. With further mass media inventions like radio and television the kind of content changed, but not the fundamental character of mass media. The second incision is the digitalisation of media. The true radical change happened after a transition phase in which content was transferred to the digital sphere more or less unchanged. Only after the introduction of what has been named “Web 2.0” the character of media changed substantially and the ground for a new kind of media was laid: media that is personalised for a specific consumer – “individual media”.

The development from individual production, to mass production, to mass individualisation is similar in other areas of consumption: From products created by artisans, to goods mass-produced in factories, to automated on-demand production of customised products. It is worth pointing out that none of these innovations has completely replaced the other ones, only the focus has shifted or is still shifting.

So far, in media, mass-individualisation is only applied as filtering existing content. Common examples of this are personalised search results, filtered feeds and compilations (like Spotify’s “Mix of the Week” or Netflix’ recommendations), all based on the consumer’s previous behaviour. In contrast to this, individual media is the next step, with content being created to fit the recipients personal needs and wishes.

Some experiments were already made into this direction, for example by using biofeedback technology to adapt content to the user’s emotional response. “Neuro Fiction” presents a story that changes based on the reader’s brainwaves which are measured as she is reading [11]. A similar approach is taken by “Nevermind”, a game using heart rate sensors and facial expression recognition to determine the player’s stress level and adjust the game accordingly [8]. It is not hard to imagine further areas to which this development towards customised content could soon be applied to: Even today new songs are being created in the style of an artist [7] or a genre [6] by or with the support of AI, and existing songs are being transferred to a new style [16]. While these experiments are still in their early stages, the implications are clear: At some point, rather than Spotify recommending existing songs for you to listen to, it might create new songs specifically for you.

However, rather than considering entertainment we put our focus on problem oriented applications of these new technologies in order to improve the information comprehension capabilities of humans.

Approaches to Fit the Information Value to the User's Needs

As mentioned in the introduction we see "Information Overload" and "Unstructured Information" as the main challenges obstructing improved comprehension of information.

Information Overload: With the amount of accessible information steadily increasing, human attention has become the scarcest resource [13]. Too much available information together with limited time means it is hard to select what to focus on – both in terms of texts available and the length of the texts themselves. Herbert Simon said: "If computers are to be helpful to us at all, it must not be in producing more information – we already have enough to occupy us from dawn to dusk – but to help us to attend to the information that is the most useful or interesting or, by whatever criteria you use, the most valuable information" [5]. With the scarcest resource being human attention, one potential way for more rapid comprehension of information is to filter and summarise content.

Unstructured Information: While hyper-media has improved the way information is connected, to goal of having "the right information" always within close reach is still not achieved (the popularity and necessity of search engines demonstrates this). A problem of written language is that the author has to make many assumptions about the reader's existing knowledge and intention. This results in texts that cannot fulfil the information necessities of every single reader. Plato famously criticised written language for this in his "Phaedrus" dialog [10]: For him, spoken language is superior, because the speaker can (and has to) adapt her words to the listener – the information becomes "alive". A teacher can ask and answer questions, react to critique and go into more detail – she can add additional information if necessary. A written document cannot do this, and so the reader is left on their own, with an inanimate piece of information where getting additional information requires extra steps to be taken by the user. Because most information is unstructured and unconnected, gaining a better comprehension of information is difficult.

Subsequently we present two conceptual applications, which try to solve these problems by utilising AI as it will likely be available in the future. Their respective goals are helping comprehending information faster by *reducing* the content, and comprehending information better by *extending* the content. The first application, *Abstraction Slider* is designed to enable the user to select how much time they want to spend on the text. They can get a quick overview over the contents of a text (and then decide what to spend time on) or read the full version, depending on their needs. This can help the user comprehend information faster. The second application, *Content Augmentation*, is designed to adapt the text to the user's information needs by extending the existing content with additional information. This can help the user to gain a better, deeper comprehension of the information.

We imagine the applications to not just modify the text based on its content, but also based on what is known about the reader – his existing knowledge, intention and context. This means truly individualised results and content exactly matching the user's need of information value being created.

Both applications try to enable the user to see the text not as a static piece of information, but as one which can be

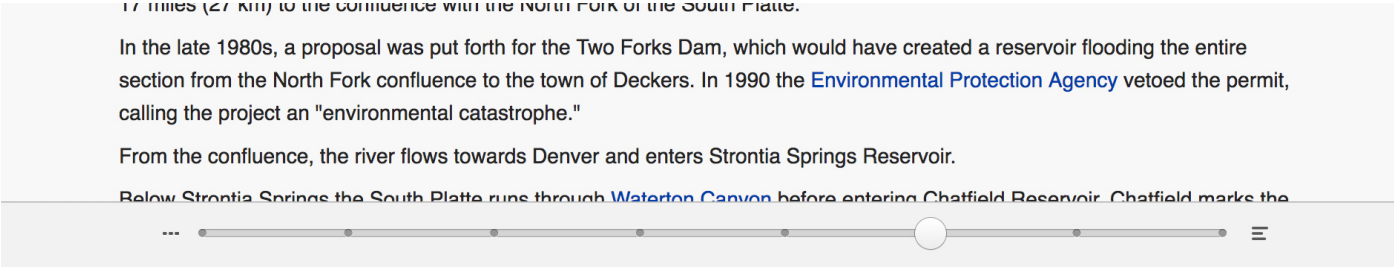

Figure 2: The Abstraction Slider Interface

modified to her will and is changed based on her specific needs and context. They both manipulate text, as it is the most common, flexible, scannable and versatile medium for information transfer. They should work with any text that is given as an input. In the following they are displayed as additions to a webbrowser user interface, because the browser is were a majority of digital written information is being consumed and available in an open format.

Application I: Abstraction Slider
In order to help users gain a more rapid comprehension of a text, the first application enables the user to control the amount of information, by choosing a level of summarisation, ranging from the original text down to keywords.

This can help users to choose which texts to spend more time on. In scientific writing this has long been practised by adding abstracts to the top of publications, so that readers can decide whether it is worthwhile their attention for their current problem. Our application tries to offer this benefit for any text and for any user. It can also help users gain a faster, though not as detailed understanding of the contents of a text, if this is what they require, eg. because of limited time. This application assumes the “capability to reorganize, modify and merge information” [9] of an AI. Based on an abstractive summarisation method any text can be effectively summarised to any stage of abstraction. With the Abstraction Slider users retain control over the level of summarisation to prevent the feeling of being patronized by the AI.

Structured interviews with a Wizard-of-Oz-Prototype were conducted, where the text was summarised manually beforehand, but the users where led to believe it was happening automatically. In the beginning the users were asked to repeat the key aspects of a text after being given two minutes to read it. After introducing them to the Abstraction Slider, the same task was repeated with a different text of similar length, but this time they were allowed to use the Abstraction Slider and had only one minute to read the text. The tests revealed that all users gained a better understanding of the core points with the Abstraction Slider, compared to reading a text of similar length and complexity without it. Most users felt comfortable with the idea of text being dynamically changed by an AI, as they could always adjust the level and compare it to the original. The need for structure became apparent: Users preferred stages where the text had been shortened, but headings had been left in

place. Additionally they expected the AI to not only sum up the texts, but simplify the language and the content as well. The interviews showed that in order to obtain a high quality of summarisation an AI has to focus on four text characteristics: structure, content, language and type.

Structure: The observations confirmed that a clear hierarchical formatting is important in order to consume information more effectively. Even if the amount of information stayed the same, all candidates preferred an abstraction level with a clear visual structure.

Content: For a natural output the text has to be rewritten considering the meaning of the information. With an increasing abstraction level redundant and irrelevant information should be eliminated and detailed information condensed.

Language: The concept of translating detailed information to more generalised information should be used across all levels of abstraction. Adjusting the abstraction level should not only affect entire paragraphs, but words and phrases also need to adapt their informative content. To increase the readability the language as well as the diction has to alter: specific information like “1929” should be translated to a generalized one like “at the beginning of the 20th century” depending on the level of abstraction.

Type: While analysing the four common text types – descriptive, narrative, expository and argumentative – user tests showed the different impacts and challenges of summarisation. Expository text driven by facts and data can be easily summarised with the help of an AI. The summarisation of a descriptive or narrative text with a strong storytelling aspect is more challenging, even though the AI can summarise the plot, the entertaining character of a story might be missing in a higher level of abstraction. An argumentative text is challenging as small changes or misinterpretations can completely distort the creator’s point of view.

The essential question is: When does summarisation make sense? By helping users gain control over the amount of information they consume the application can help them deal with “Information Overflow” and focus on relevant information.

Application II: Content Augmentation
The second application, called “Content Augmentation”, tries to enable users to get into an interactive dialog with the text, thereby gaining a deeper understanding of the information presented. Technically the application can be seen as a human-friendly interface for an intelligent agent [15] that collects additional information for the text at hand.

One of the biggest problems of written language is its static state: An author has to make many assumptions about the intended readers. Restricted by static media texts are addressing a “mean expected reader” and are created without being able to react to different readers with specific necessities. Through the “Content Augmentation” interface text is treated as a dynamic and reactive media. The AI can be imagined as a knowledgable third party, that guides the users and tries to help them comprehend complex information better.

Based on the hypothesis: “Users gain a better comprehension when specific sections of text can be augmented”, texts of different types were analysed. As a result we categorised the information missing for better comprehension into four main kinds: explanatory, extending, opinion-forming and proving information. In the application the user can interact with highlighted parts of the text via predefined operations in order to be presented with additional information of above mentioned kinds. The interface visually blends

Figure 3: The Content Augmentation Interface (after the text is expanded)

on top of the existing text. If the AI can offer additional value for a paragraph, one or more of the following operations are available:

Tell me more: This operation replaces the highlighted section with a more detailed version. When the user is interested in a statement, the AI can elaborate the existing text and rewrite the current section. Additional information can be added, for example explanations about events, persons or organisations mentioned.

Explain: When writing a text the author has to assume topic related knowledge. With this operation the AI provides a simplified version of the paragraph, by adjusting the register, diction and minimizing specialist knowledge. New information should only be added to make the current statement more clear and to improve comprehension.

Validate: Within written text references to other publications often occur. This operation is used to reveal the referenced source of a statement spotted by the AI. Therefore the semantic context of the highlighted paragraph has to be identified and suitable references be found.

Agree/Disagree: During out concept phase the possibility of opinions being expressed more clearly (or being disputed) by the AI became apparent. Based on the users opinion the text might change to enable a more fruitful discourse, providing them with supporting and opposing views.

The application tries to facilitate a better and easier comprehension by allowing users to not accept the static form of a text, but request explanations, get additional information and validate sections. With the help of these interactive operations the text becomes alive and shifts from a static to a dynamic medium. It is obvious that this application cannot replace a direct discussion with a "teacher", but written language persisted despite Plato's critique, because of its obvious benefits. With our concepts we hope to lessen the impact of its drawbacks to enable users to gain a better comprehension of written information.

Potential Impacts of Individualised Content

Automatic content creation and summarisation are technologies already in use. So are personalised feeds, recommendations and predictions for the user's behaviour. Combining these technologies into what we call individual media is inevitable to happen, even if the AI needs to be more powerful than current AI. It is important to be aware of the implications that this technology will have for single users, as well as society as a whole.

Some of these consequences of the existing personalisation technologies are already being discussed: "Filter Bubbles" have emerged, which lead to users finding themselves in echo chambers, not being confronted with any information opposing their views. This ultimately widens the gap between different groups in society. Similar technology is also used by targeted advertising, were tailored content aims to influence the user's decisions – usually without their knowledge. These developments are the forerunners of what could happen if media is being created for the individual without any reflection. When AI directly manipulates the content that is being consumed, its power will be massive. If the AI is trained to reduce or extend the content skewed towards certain political views, or impelled by economic incentives rather than by the user's best interest the implications might be severe. Deciding which parts of a text are more important than others and which informations need to be connected is not an easy task – not for humans today and not for AI in the future. A "neutral" manipulation of a text probably does not exist. There is also the issue of humans trusting machines too much: studies have shown that "Overtrust" in machines is something to be reckoned with [12]. This implies that those creating these technologies need to be very careful in considering who they support and what their motives are. The responsibility of those developing and training the AI is bigger than ever before.

With the concept of individual media, where content is dynamic and individualised, the personal comprehension of a user might be improved. But we should not lose sight of the challenge to improve the shared comprehension of information. A world were individual media is dominant might be a world were no two people see the same thing when they read a text. The implications of this will need to be evaluated with great caution: Maybe it's just similar to two people having a conversation, each with a different teacher, about the same text – they still share the same core information, but the perspective gained can be vastly different. But it is also possible that this leads to great uncertainty over what can be seen as common knowledge and a greater segmentation of society, where the "Filter Bubble" is expanded to all areas of information consumption. When the access to knowledge becomes individualised – not just in the way it is now, but in a far greater sense – the difference between a fact and an opinion (even if held by an AI) is blurring.

The enormous consequences of the introduction of individual media should serve as a call to action for engineers and designers. They should start focussing now on how to shape this technology in a way that is the most beneficial to society. Engineers should help developing AI that is controllable and as transparent as possible. If an AI cannot be held accountable for its actions, it is the liability of the developer to make sure it is safe and its decisions are verifiable. We believe it is important that designers put their focus on how to make the decisions made by the AI transparent to the users and help them retain power. As AI starts to influence day-to-day life more and more, users need to be made aware of its role and have to be sensitised to its presence. It is vital that systems are designed to encourage humans not to lose their critical view and to continue reflecting on the machine's decisions. Especially as AI is being applied

to such crucial domains as how we comprehend information, potential for misuse and manipulation is created.

It is the responsibility of engineers, designers and creators in general to not underestimate this change in the nature of media and to put the needs of humans in center of the development of new applications. These duties begin now, as the transition phase to a world coined by intelligent machines accelerates. Individual media has great potential to help us “Augmenting Human Intellect”. It can provide us with a long term gain, if it is used to enhance our ability to comprehend information and help us approach complex problems, that we would not be able to derive solutions to otherwise.

Conclusion

Individual media is the next step in the progression of media. Its ramifications for how humans consume information are big. If used right it can help improve the way humans comprehend information, resulting in an improved ability to find solutions for complex problems.

The “Abstraction Slider” and the “Content Augmentation” fit the information value to the userâĂŹs need, by summarising and augmenting content as necessary. This helps users to only read the information required for their goal, while enabling them to gain insights not provided by the original text.

To ensure the positive use of individual media technology, engineers need to focus on developing the right AI for these purposes, while being aware of their specific dangers. Designers need to start now to find ways to help non-technical users to understand the decisions of the AI and how they affect what they see, while enabling them to use it to its full potential.

The two applications presented in this paper aspire to be examples of this. They present simple interfaces, that minimise the difficulty of using powerful AI. But they also allow users to compare the modified content to the original untouched version, so that they can check and examine which changes the AI is introducing.

References

[1] Ilze Birzniece. 2011. Artificial Intelligence in Knowledge Management: Overview and Trends. *Scientific Journal of Riga Technical University. Computer Sciences* 43, 1 (2011), 5–11. DOI: http://dx.doi.org/10.2478/v10143-011-0001-x

[2] Vannevar Bush. 1945. As We May Think. *Atlantic Monthly* 176 (July 1945), 101–108.

[3] L. Diao, M. Zuo, and Q. Liu. 2009. The Artificial Intelligence in Personal Knowledge Management. In *Second International Symposium on Knowledge Acquisition and Modeling, 2009. KAM ’09*, Vol. 3. 327–329. DOI: http://dx.doi.org/10.1109/KAM.2009.269

[4] Douglas C. Engelbart. 1962. Augmenting human intellect: A Conceptual Framework. (1962). http://www.citeulike.org/group/22/article/1030610

[5] G. Fischer and J. Otswald. 2001. Knowledge management: problems, promises, realities, and challenges. *IEEE Intelligent Systems* 16, 1 (Jan. 2001), 60–72. DOI: http://dx.doi.org/10.1109/5254.912386

[6] Ji-Sung Kim. 2016. Deep learning driven jazz generation. (March 2016). https://jisungk.github.io/deepjazz/

[7] SONY CSL Research Laboratory. 2016. AI makes pop music in different music styles. (Sept. 2016). http://www.flow-machines.com/ai-makes-pop-music/

[8] Adam Lobel, Marientina Gotsis, Erin Reynolds, Michael Annetta, Rutger C.M.E. Engels, and Isabela Granic. 2016. Designing and Utilizing Biofeedback Games for Emotion Regulation: The Case of Never-

mind. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. ACM, New York, NY, USA, 1945–1951. DOI: http://dx.doi.org/10.1145/2851581.2892521

[9] Ani Nenkova and Kathleen McKeown. 2011. Automatic Summarization. *Foundations and Trends in Information Retrieval* 5, 2–3 (June 2011), 103–233. DOI: http://dx.doi.org/10.1561/1500000015

[10] Plato. 1997. Phaedrus. In *Plato: Complete Works*, John M. Cooper (Ed.). 228b.

[11] Hannu Rajaniemi and Samuel Halliday. 2013. Neurofiction. (2013). http://neurofiction.net

[12] Paul Robinette, Wenchen Li, Robert Allen, Ayanna M. Howard, and Alan R. Wagner. 2016. Overtrust of Robots in Emergency Evacuation Scenarios. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction (HRI '16)*. IEEE Press, Piscataway, NJ, USA, 101–108. http://dl.acm.org/citation.cfm?id=2906831.2906851

[13] Claudia Roda. 2014. *Human attention in digital environments* (1. paperback ed. ed.). Univ. Press, S.l.

[14] Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A Neural Attention Model for Abstractive Sentence Summarization. *arXiv:1509.00685 [cs]* (Sept. 2015). http://arxiv.org/abs/1509.00685 arXiv: 1509.00685.

[15] Leon Sterling. 2006. Knowledge Management Agents. In *Encyclopedia of Knowledge Management*, David Schwartz (Ed.). 368–372.

[16] Dmitry Ulyanov and Vadim Lebedev. 2016. Audio texture synthesis and style transfer. (Dec. 2016). http://dmitryulyanov.github.io/audio-texture-synthesis-and-style-transfer/