



Exploring Communicative AI: Reflections from a Swedish Newsroom

Downloaded from: <https://research.chalmers.se>, 2025-12-04 22:41 UTC

Citation for the original published paper (version of record):

Stenbom, A., Wiggberg, M., Norlund, T. (2023). Exploring Communicative AI: Reflections from a Swedish Newsroom. *Digital Journalism*, 11(9): 1622-1640.

<http://dx.doi.org/10.1080/21670811.2021.2007781>

N.B. When citing this work, cite the original published paper.

Exploring Communicative AI: Reflections from a Swedish Newsroom

Agnes Stenbom^{a,b} , Mattias Wiggberg^b  and Tobias Norlund^c 

^aSchibsted, Stockholm, Sweden; ^bThe Department of Industrial Economics and Management, KTH Royal Institute of Technology, Stockholm, Sweden; ^cDepartment of Computer Science and Engineering, Chalmers University of Technology, Gothenburg, Sweden

ABSTRACT

This article contributes to the emerging field of research on computational journalism with a practical illustration of an attempt to utilize Machine Learning to generate Search Engine Optimized headlines in a major Swedish newsroom. By using its technical results as a springboard for reflections among internal stakeholders, the experiment serves as a catalyzing innovation revealing deliberations on computational approaches in journalism in general and communicative Artificial Intelligence (AI) in specific. The study concludes with three ideas to support decision makers involved in evaluating potential use cases for communicative AI in journalism.

KEYWORDS

Computational journalism; machine learning; natural language generation; search engine optimization; human-machine communication; communicative AI

Introduction

Digital journalism studies have been described as “*an academic field which critically explores, documents, and explains the interplay of digitization and journalism, continuity and change*” (Eldridge et al. 2019). Recent research in the field has engaged with the growth of automated practices, new journalistic routines, and contemporary issues related to Artificial Intelligence (AI), with much scholarship stemming from the evolving concept of machines performing communicative tasks in the news production chain. But what is machine communication, and why should journalism care?

Guzman and Lewis (2020) use the term *communicative AI* to explain applications such as e.g. automated-writing software or social bots that are “*designed to carry out specific tasks within the communication process that have formerly been associated with humans*”. Communicative AI enables machines to function as communicators themselves, rather than mere mediators of human communication. A prominent technological concept enabling such applications is found in the Machine Learning (ML) process of Natural Language Generation (NLG), where computer software synthesizes natural language such as text or speech. The innovation occurring when such digital technologies are introduced in newsrooms is important for media organizations to understand in order to improve their chances of long-term survival (Küng 2015). However, in studies of newsroom innovation, a dichotomy

CONTACT Agnes Stenbom  astenbom@kth.se

© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

of human-centrism and techno-centrism tends to appear as attention is guided to *either* journalists or technology (De Maeyer 2016). With this work, we strive to balance technological and social dimensions as we seek to understand an innovation process related to communicative AI in journalism. Aligned with the proposition of Lewis, Guzman, and Schmidt (2019), we seek to highlight how machine communication is not an all-or-nothing proposition and aim to take into account both communicative abilities and limitations as we approach an experiment at a major Swedish newsroom aimed at using NLG to generate organic search traffic.

Background

A rapid societal transition toward algorithmic services such as social media and search engines is transforming the institution of news media by undercutting business models, altering work routines, and providing vast information alternatives (Broussard et al. 2019). As such, media organizations around the world are experimenting with new ways to serve audiences and sustain business (Duffy and Ang 2019). An associated challenge to this is that of attaining the attention of news readers, and a key concept to consider in this regard is that of Search Engine Optimization (SEO) – a term covering a range of activities with the goal of increasing organic traffic referrals to websites.

Applied in news media, SEO has been described as incorporating techniques aiming for high rankings on search engine results pages into journalistic processes, leading content to be written to adhere to search engines' (rather than humans') definition of relevance (Dick 2011). Ranking at the top on result pages of major platforms such as Google can be critical for news publications, as the position and frequency of a site's appearance influences the traffic it receives (Giomelakis, Karypidou, and Veglis 2019). Together with social media, search engines are one of the primary ways digital readers discover news (Newman 2011).

In this article, we study an experiment at the Swedish publisher Aftonbladet with the goal of innovating the way headlines are optimized for search engines through employing communicative AI. We investigate two questions:

- Can NLG can be utilized to generate SEO headlines of sufficiently high quality for Aftonbladet?
- What are the stakeholders' reflections on the experiment in specific, and communicative AI in the newsroom more broadly?

Inspired by recent calls for research in the Human–Machine Communication (HMC) domain (see e.g., Lewis et al. (2019), we focus specifically on how the stakeholders come to see themselves in relation to the communicating machine, and the machine in relation to them. While disclosing the technical setup of the experiment, the key findings presented in this article therefore concern the stakeholders' reflections.

Related Work

The capabilities of computers are increasingly utilized in journalism and the boundaries that have historically surrounded the field are changing, with audiences,

algorithms, and analytics now serving as integral parts of modern news work (Duffy and Ang 2019). Algorithms today can prioritize, classify, and filter information, thereby making their way into several stages of the journalistic process (Anderson 2013; Carlson 2015). There is no one term fully capturing this notion of utilizing data and computational approaches during journalistic activities. Various terms such as *automated* or *robotic journalism* (Latzer et al. 2014), *data journalism* (Coddington 2015) or *algorithmic journalism* (Dörr 2016) have been used, and the broad term *computational journalism* has been used to encompass many of them. Cohen et al. (2011) define computational journalism as “*the combination of algorithms, data, and knowledge from the social sciences to supplement the accountability function of journalism*”. Diakopoulos (2019b) explains the term as “*information and knowledge production with, by, and about algorithms that embraces journalistic values*”. Used in a variety of newsrooms around the world (including the Washington Post, Bloomberg, and the LA Times; Tandoc et al. 2020), computational journalism is presented as a means to provide newsrooms with speedy and scalable supplies of narratives as media organizations’ financial resources decrease (Graefe 2016).

While media organizations often define themselves through human content-creation activities, the modern media industry is largely a technology industry. Now, technology is a contributor to – not just enabler of – the generation of media content, leading the two to co-exist in a symbiotic relationship (Küng 2013). New journalistic practices are emerging as programming and journalism merge (Coddington 2015), and as noted by Deuze and Witschge (2018), “*what journalism is and what journalists do cannot be meaningfully separated from their material context*”. Elegantly summarizing what this means for digital journalism studies, Lewis and Westlund (2015) suggest framing studies in terms of a collaboration of human actors and technological actants, as all may be “*intertwined in the activities that constitute cross-media news work*”.

These developments bring new issues for news organizations, including but not limited to the intricate challenges associated with blending social and occupational worlds through increasingly blurred boundaries (Lewis and Usher 2016). Parallel to hybrid approaches – here understood as processes combining algorithms, automation, and people to conduct news work (Diakopoulos 2019a) – growing in number, there is also concerns of human editorial workers becoming abundant (Broussard 2018). In a global survey on newsroom attitudes to AI, fear of job replacement was noted among the most prominent issues (Beckett 2019). The professional group of reporters and correspondents has been suggested to face an 11% automation risk (Frey and Osborne 2017). While the reliability of such predictions has been questioned, among others by Nedelkoska and Quintini (2018), suggesting that there is a tendency towards automation in the news media industry is not a controversial idea. Facing budget constraints, many news organizations are looking to automation to streamline manpower, cut costs, and improve efficiency (Graefe 2016; Wu et al. 2019; Tandoc et al. 2020).

In relation to these developments, Wu et al. (2019) illustrate how technologists have earned an understanding of the current economic situation of journalism, which has allowed them to position automation as a potential remedy; a tool which will help news organizations meet audience demands as resources decline, and free up journalists to work on more in-depth reports. The editor of news production systems at

Reuters has shared a similar idea, noting how the goal of their experimentation with automated forms of journalism is “*not to take anyone’s job*” but to identify ways in which humans and machines can collaborate in news production (Bilton 2018)

The implications of the shift towards automation in the newsroom have awoken academic, societal, political, and commercial debates concerning the role of media organizations and the process of creating and distributing media content through traditionally human activities of creativity and expressions (Latzer et al. 2014). The rapid digitalization and changed media landscape of the past two decades has introduced new challenges, including issues such as filter bubbles, fake news, fragmentation, and click-driven journalism (Syvertsen et al. 2019). Novel questions surrounding the societal responsibility of news organizations are arising. Now, the issue at hand is not necessarily whether a given computational approach is *possible*, but whether it is *desirable* – both in regards to the journalism produced and the process behind it.

Theoretical Framework

With new opportunities and challenges brought to journalism through computational innovations and automation, updated theoretical frameworks are needed for understanding the ongoing transformation. One such which finds particular relevance in regards to the type of technological implementation studied in this research is HMC. HMC is an emerging area within communication research studying the “*creation of meaning among humans and machines*” (Guzman 2018). By providing theoretical frames for conceptualizing the meaning of being a communicator, HMC challenges the traditional anthropocentric assumption of humans being the originators (and machines only mediators) of communication by provoking inquiry about what happens when machine systems take on roles previously held by humans (Lewis et al. 2019). As outlined by Guzman (2018), the goal of HMC is to develop theories of communication for human-to-machine communication contexts and better understand the larger social and cultural implications of devices and programs taking on communicative roles. As suggested by Guzman and Lewis (2020), the focus of HMC research can be on functional, relational or even metaphysical aspects.

The HMC framework is useful for researchers concerned with the transformation of journalism as it assists our understanding of two key areas of inquiry: how machines take on roles as communicators in journalism, and how humans understand and perceive them in such. The rationale for employing HMC in this study is found in a quest for understanding the latter. Using HMC, we are able to address how stakeholders in a newsroom experiment reflect on the idea of machines as communicators in journalism and theorize based on our findings.

Research Context

This research concerns an experiment occurring at the major Swedish publisher Aftonbladet. With up to 50 million daily total page views (Lund and Kårhus 2020) in a country of 10.4 million (OECD (2021)), Aftonbladet is a national digital news destination. Aftonbladet is part of the publicly traded company Schibsted, a Nordic group of

Table 1. Examples of journalistic and SEO headlines from Aftonbladet, September 2020 (the English versions are translated by the authors).

Journalistic headline (displayed on Aftonbladet)	SEO headline (displayed on Google)
Ruth Bader Ginsburg död i cancer – politisk strid direkt (Ruth Bader Ginsburg dead in cancer – political battle directly) Sex cancersymtom du ska ta på allvar (Six cancer symptoms you should take seriously)	Ikoniska HD-domaren Ruth Bader Ginsburg är död – blev 87 år (Iconic Supreme Court judge Ruth Bader Ginsburg is dead – became 87 years) Sex symtom på cancer du bör söka vård för (Six symptoms of cancer you should seek care for)

digital consumer brands with approximately five thousand employees and forty different brands. There are two additional news organizations within the Swedish Schibsted-ecosystem: the legacy publisher Svenska Dagbladet and the news aggregator service Omni.

Aftonbladet is a predominantly ad-funded news destination, and the amount of ads sold directly relate to the amount of readers (the traffic) of the legacy publisher's site. Efforts for increased traffic are thus business critical, and there is a dedicated department working to promote SEO practices across the newsroom. Their efforts in the domain are in large part focused on tags and headlines. Journalists are asked to provide two headlines for each new article they write: a *journalistic headline* displayed on the Aftonbladet site, and an *SEO headline* displayed on the Google search engine results page. While journalistic headlines at Aftonbladet leverage creativity, cultural associations, and elements of surprise, SEO headlines are descriptive and contain words people typically use to search for content (Table 1).

In 2019, the SEO-team learned that this practice of dual headline production was perceived as tedious among journalists and were therefore looking for ways to innovate their practices through utilizing new technologies. It is in this context that the below described experiment was carried out.

Method

Our inquiry into the SEO-experiment occurring at Aftonbladet was done through a sequential research design in which a technical experiment was carried out and its result later informed reflections among its stakeholders, captured through semi-structured interviews. This was done through an action research approach.

Getting Close to the Exploration through an Action Approach

Getting empirical insight into innovation processes of a media organization is an intricate challenge. Wagemans and Witschge (2019) discuss how the *“iterative nature of innovation processes in the current media landscape makes it difficult to study these phenomena in a linear fashion”*. Accessing organizations, understanding what is happening, and earning confidence to be involved in innovation processes with sensitive commercial results are all areas that make the phenomenon difficult to study. To remedy some of these challenges, leveraging elements of action research can be highly useful. Action research has been described as *“evaluating your practice to check whether it is as good as you would like it to be, identifying any areas that you feel need*

improving, and finding ways to improve them" (McNiff 2016). Action research approaches have been established in journalism studies with noteworthy illustrations related to innovation and practice improvement (Posetti 2018; Grubenmann 2016; Wagemans and Witschge 2019).

In this study, two of the authors participated in the below described experiment as employees at Schibsted with the typical action research aim of "[...] *bringing theory and practice closer together*" (Karlsson and Sjøvaag 2018). This allowed for an action inquiry approach where the researchers had deep involvement in the process studied (Becker and Huselid 2006). Action inquiry (which is to be understood as part of the broader domain of action research) enables scholars to generate timely action (Torbert and Taylor 2008). The approach has been described as "*a way of simultaneously conducting action and inquiry as a disciplined leadership practice that increased the wider effectiveness of our actions*" (Torbert 2004). More broadly speaking, the concepts of action inquiry assert the proposition of interpretive social science, suggesting how the social world must be understood from within rather than explained from without (Hollis 2002). Being on the inside of the studied phenomena – in this case, the SEO-experiment – thus enables researchers to deeply inquire about it.

The action inquiry approach enabled privileged access to data informing this study. Another example of privileged data access from the domain of digital journalism studies can be found in Jones and Jones (2019b) inquiry into innovation processes at the BBC, where they themselves were also employed and thereby gained deep insight into the transformation they were studying. Inspired by their approach, this work leveraged already established relationships with decision makers in the news organization studied to enable insight into processes which informed this research.

Our research setup can be understood as occurring in two primary phases, with action inquiry enabling the shift from the first to the latter. First, the authors were involved in the process of experimenting with the hypothesis of machine-generating SEO-headlines sparked at the Aftonbladet in 2019. The authors then held roles as Data Scientist and Product Manager employed at Schibsted, and worked together with the SEO team of Aftonbladet in the experiment. They were thereby deeply involved in the setup of the experiment, but not the evaluation of its performance (please find more details about the experiment and evaluation process below). As the experiment finished, the authors stepped out of the operational process and leveraged their established connection with stakeholders in the experiment to learn about their reflections. This second stage thus constituted a process of "zooming out" from the operational experiment to understand and position the work in a theoretical context.

The Experiment

NLG is frequently used for communicative tasks such as translation (Bahdanau, Cho, and Bengio 2015) and summarizing text (Allahyari et al. 2017). Significant progress has been made in the field in recent years, with notable developments including modeling improvements (Bahdanau et al. 2015; Vaswani et al. 2017) and increased amounts of data and computing power accessible for training (Radford Alec et al. 2019; Raffel et al. 2019; Brown et al. 2020). After learning about these developments, staff from the SEO team at Aftonbladet approached the machine learning team of Schibsted

with a hypothesis on using NLG to generate SEO-headlines. Their idea included 1) generating SEO headlines for *historical* articles to gain value from already performed journalistic work by driving traffic to the online news site, and 2) generating SEO headline suggestions to journalists in order to drive traffic to *future* articles. After discussing the general idea, the two teams set up a collaborative experiment to explore its potential, running from the fall of 2019 into the early spring of 2020.

In this experiment, a novel hybrid model architecture was developed based on the Pointer Generator (See et al. 2017) and Transformer (Vaswani et al. 2017) models. Two different datasets were utilized to train the model: *Schibsted-articles* and *Aftonbladet-seo-articles*. *Schibsted-articles* contained 2,967,497 articles from Schibsted's Swedish publishers Aftonbladet, Svenska Dagbladet and Omni with each article including a headline, body and metadata (tags). *Aftonbladet-seo-articles* contained 83,317 articles from Aftonbladet which, in addition to the headline, body and metadata, also included a human-written SEO headline. Using these two datasets as training data, the model was first taught to generate generic journalistic headlines from the full Schibsted-dataset, and then fine-tune its performance on the Aftonbladet-dataset to generate SEO-headlines. Further background and technical details about the training process are described in [Appendix A.1](#).

Evaluation Process

After training the final model, the quality of its generated headlines was to be evaluated by human stakeholders. The final test of the model was conducted on 62 randomly sampled articles published on Aftonbladet via the Swedish news agency Tidningarnas Telegrambyrå (TT). This was motivated by an ambition to create a comparable dataset; the TT-articles were considered to keep to a similar style and structure, and the risk of individual journalists' writing styles influencing the model's performance would thereby be minimized.

The model generated five SEO headlines for each of the 62 articles, presented according to model-assumed relevance. A dedicated SEO Expert from Aftonbladet then reviewed them, starting with the model-assumed top alternative and then considering whether any of the top 1–5 generated headlines was of sufficient quality for *theoretical* publication. The evaluation was done through an online sheet set up by the ML team in which the SEO Expert had two possible boxes to check, one asking if the assumed top alternative was approved and the other asking if any of the headlines were.

The SEO and ML teams collectively set the goal of 80% accuracy, determining that eight out of 10 suggested headlines must be approved by the SEO Expert in order for the experiment to continue into wider testing and implementation. As there is no one definition of what constitutes a *correct* SEO-headline, it was the job of the SEO Expert to use their professional expertise to evaluate whether the machine had indeed learned to generate SEO-headlines.

Stakeholder Interviews

Leveraging a common method used in digital journalism studies (see e.g., Thurman et al. (2017)), semi-structured interviews were conducted with key stakeholders. The

interviews were conducted to gather diverse perspectives on the experiment and enable deeper understanding of the participating stakeholders' attitudes and reflections on it.

The informants were purposefully limited to individuals involved in the project, ultimately yielding three candidates for interviews: an SEO Expert, an ML Expert and a Strategist. The SEO perspective was deemed relevant to understand the domain of SEO and challenges related to inadequate such. The ML perspective was deemed relevant to understand technical challenges associated with utilizing ML to generate natural language. Finally, the Strategy perspective was deemed relevant to understand the broader context of decisions related to the experiment. As the Strategy team was not part of the practical experiment in the same capacity as the SEO and ML teams, but rather functioned as a supervisory team, this interview aimed at understanding overarching considerations.

A general interview guide was created for consistency in the interviews, yet as the informants represented different domains, additional questions were asked related to the respective informants' domain of expertise. The interview guide aimed to maintain reliability and enable valid analysis of differing perspectives. It also served to mitigate biased interview approaches, which was especially important as the interviewer was involved in the experiment. The interviews were conducted in a mix of Swedish and Norwegian (also known as Scandinavian). Before the interview, informants were briefed on the purpose of the study and the relevance of their perspective.

One researcher conducted the interview over video conferencing, and the interview was recorded and subsequently transcribed. Each informant was asked to review the gathered information, add comments or correct potential misunderstandings before confirming the accuracy of the data.

The interview data were analyzed with a grounded theory approach (Vollstedt and Rezat, 2019), a common method in digital journalism (Steensen and Ahva, 2015), with a goal of uncovering challenges associated with the experiment. In this process, the focus was primarily but not exclusively on the concept of machines functioning as communicators in journalism

Findings

As presented, the research design of this study consisted of two phases: a technical experiment and subsequent stakeholder reflections on it. In the following paragraphs, we present the findings accordingly.

Machine Learning-Generated SEO Headlines

The majority of articles were provided appropriate SEO headlines through the machine generation, with successful examples including adding events (such as *Tour De Ski*) and first names (*Charlotte Kalla*). However, the model failed to consistently generate accurate results. This led some of the ML-generated headlines to diverge from the original meaning of the article. Unsuccessful examples include removing important

Table 2. Examples of ML-generated SEO headlines (English versions translated by authors).

Original headline	ML generated SEO headline (1–5)	Any of the	
		First headline: yes/no	headlines 1–5: yes/no
Kalla bröt staven: "Blev irriterad" (Kalla broke the pole: "Got irritated")	Charlotte Kalla bröt staven i Tour de Ski (Charlotte Kalla broke the pole in Tour de Ski) Charlotte Kalla bröt staven: "Blev irriterad" (Charlotte Kalla broke the pole in Tour de Ski: "Got irritated") Charlotte Kalla bröt staven (Charlotte Kalla broke the pole) Charlotte Kalla bröt Tour de Ski (Charlotte Kalla broke Tour de Ski) Charlotte Kalla bröt stav i Tour de Ski (Charlotte Kalla broke pole in Tour de Ski)	Yes	Yes
Kurz tillbaka Österrike vid makten i Österrike (Kurz back in power in Austria)	Sebastian Kurz tillbaka i Österrike (Sebastian Kurz back in Austria) Kurz Kurz tillbaka i Österrike (Kurz Kurz back in Austria) Kurz VÖ tillbaka i Österrike (Kurz VÖ back in Austria) Kurz VÖ tillbaka vid makten i Österrike (Kurz VÖ back in power in Austria) Kurz Kurz tillbaka vid makten i Österrike (Kurz Kurz back in power in Austria)	No	No

information (Kurz back *in power* in Austria), incorrect grammar such as repeating words, or providing incorrect or vague information about complex events. (Table 2)

When reviewing the model-assumed top 1–5 generated headlines, the SEO Expert approved 61%. The headlines assumed most relevant (i.e., top 1), however, were accepted by the SEO Expert in only 45% of the cases. The set goal of 80% acceptance was hence not reached.

Stakeholder Reflections

Discussing these findings with stakeholders in the experiment, three key themes emerged in the interviews. Some related directly to the carried-out experiment and others to communicative AI more broadly, including considerations about low-quality technical solutions applied in newsrooms, cultural challenges emerging from interactions between technology and editorial staff, and challenges related to fulfilling the societal responsibilities of news organizations when using communicative AI.

Technical Limitations

The experiment was partly discussed with a focus on technical limitations. These focused primarily on the limited capabilities for contextual understanding and the challenges of introducing insufficiently well-functioning technology to an already pressured workforce.

Lack of Contextual Understanding. The limitations of NLG were discussed in all interviews. The results of the SEO-experiment were framed as the consequence of the fundamental limitation that ML systems are currently unable to gain semantic understanding in the same capacity as human beings. As the ML Expert put it, "[Machine Learning] is *mechanical calculations*. We are not talking about any kind of

real understanding, or real thinking". For the SEO headline generation, this limitation was illustrated through the model's inability to understand the article for which it was generating a headline.

Organizational Disbelief. Related to the technical limitations of the model, the SEO Expert stressed the importance of not implementing computational solutions of insufficient quality in newsrooms, emphasizing the need for trust. *"I think that if you are to implement something like this, it has to be good from the start. Otherwise the trust in [computational approaches] will be even worse than it is [today]."*

The rationale for not introducing low quality communicative AI solutions was the trend towards generalization of the journalist role, where low quality computational solutions were perceived as risking adding stress and confusion in an already pressured work environment. The SEO Expert described how the role of journalists has transformed, noting that *"[newsrooms] used to have editors who sat and read all the texts and edited them to perfection"*, adding that the pressure on the individual journalist is much greater today than it used to be.

Framing Computational Journalism

All informants brought forward organizational challenges related to computational journalism. New themes emerged, including how to position efforts leveraging communicative AI as non-threatening towards journalists, as well as ensuring reasonable expectations.

From Threat to Everyday Relief. A prevalent theme brought forward by the informants concerned how the experiment might threaten those which it sought to unburden (journalists who were writing dual headlines for their article). The Strategist formulated how *"in all the different industries where something has been automated, there has been a sense of taking away someone's job, or professional expertise"*. Related were thoughts regarding the future need for human journalists. Posing as a journalist, the Strategist illustrated an assumed worry: *"Okay, now we'll be automated, will I not be needed?"*

Various strategies for minimizing this assumed threat were presented. The dominant argument stressed how computational approaches simplify journalists' work through performing monotonous tasks (thus lessening their workload). The SEO Expert illustrated how *"[communicative AI] is just great help for journalists. They get out of writing the kind of things a human being doesn't have to."* The Strategist suggested focusing on how things become easier through computational approaches. *"I think you need to land in the average man or woman's daily tasks [...] So, focus on the positive things."*

Positioning computational approaches as a means for journalists to focus on quality above quantity was another recurring theme. The Strategist noted how *"instead of sitting and monotonously making a site or writing a push story that lives for three hours and then goes away forever, you can instead use your energy on refining something else."* The informants stressed the need for human skills to do investigative reporting and identifying noteworthy news stories, focusing on how such craftsmanship may have long-lasting impacts on people in society. The importance of such work was further

emphasized through putting it in contrast to monotonous tasks where communicative AI could be utilized. “[We should put] *more [human] focus on things that drive us forward, not just reporting on what is happening. The latter can be done automatically to a further extent*” (the Strategist).

The informants repeatedly discussed the professional motivation of journalists and the possibility to emphasize what were believed to be more sought after tasks. The Strategist considered how journalists take employment at Aftonbladet because they want to “*make a difference*” and create widespread journalism, and how they could be motivated to work with computational tools if such were positioned as increasing the reach and impact of their human-created journalism. “*Regardless what [the benefit] is, focus should be on what fundamentally gets better.*” (the Strategist)

Expectations Management. All informants stressed the need to set achievable expectations. As communicated by the ML Expert, success is relative and expectations need to be realistic. “*Had we expected ten percent [accuracy] and gotten fifty [percent accuracy], we might have been overwhelmed*” (the ML Expert).

The hierarchy between journalists (humans) and computational approaches (machines) was another dominant theme, where machines were suggested as subordinates to humans. Building upon this idea, the ML Expert proposed the use of metaphors as means for setting achievable expectations: “*You could perhaps use a kitchen metaphor. Instead of a machine giving you an omelet, you get some egg mix with a bit of shells and an onion cut into too large pieces. It will be less work for you to cook the omelet, but you still need to take care of it. [...] It should be very clear that there is action, work, or revision expected from the humans in charge of the machines, rather than their output being presented as something done or readily cooked.*”

The ML Expert repeatedly noted that “*you need to have humans in the loop when using [communicative] AI in journalism. It’s is an incredible help in some contexts, but it’s not magic.*” The Strategist brought forward a related idea, sharing how other computational efforts at Aftonbladet have benefited from newsroom involvement throughout the project(s). However, this idea of participatory development did not go undisputed. The SEO Expert believed that editorial staff want as limited involvement in technical work as possible. The SEO Expert did note, however, that it is important that collaborating parties developing tools for the newsroom have an understanding of each other’s work, and suggested that future projects would benefit from tighter collaboration. “*Then of course it is impossible to get it to one hundred percent, because I will not one hundred percent understand what [the ML team] do and they may not one hundred percent understand what I do.*” (the SEO Expert).

Fulfilling the Publisher Mission

The final theme identified perceptions of struggles to maintain and safeguard the social mission of news publishers when using communicative AI technologies.

Vague Public Value Vocabulary. Throughout the interviews and in various ways, the respondents all reaffirmed the civic responsibility of Aftonbladet and Schibsted. Terms like *social mission, democratic mission, important tasks* and *making a difference* were all

used without further definition in the context of the importance of high-quality news media. The Strategist formulated the importance of utilizing communicative AI in trustworthy ways: *“In a ‘post truth society’ I think it is incredibly important with news media that actually are accurate and take responsibility. Given our social mission, I think it needs to be very clear that you can trust us.”*

The Need for Accuracy. All respondents agreed on the need for factually accurate news, but there were some discussion related to niche editorial verticals such as food or travel. As noted by the Strategist, *“I don’t put as high demands on recipes being accurate compared to news being accurate”*. This idea of there being more room for mistakes in non-news verticals was also extended to SEO, with the informants noting that the behaviors associated with search engines could be forgiving. The ML Expert noted how *“After a while, you may have become more accustomed to things being said online that are not true”* and questioned if factually inaccurate SEO headlines would have much negative impact on the associated brands. *“I’m not sure if this particular case would have had any significant negative consequences [...] I think it is dependent on how many mistakes it [the ML model] would have generated, how serious they were, and whether it was the mistake itself that made the article rank high on Google”* (the ML Expert).

On the topic of the need for accuracy, the informants all commented on what type of organizations they perceived Aftonbladet and Schibsted to be. Describing the organizational identity of Aftonbladet, the SEO Expert noted how *“yes, we are a digital company, but we are still a newspaper to a certain degree”*. The Strategist and ML expert also noted how there are differences between publishers and other content creators, discussing themes such as Aftonbladet’s legal responsibilities and their genuine care about what is published. This was contrasted to the behaviors of e.g., content mills putting high quantities of content online. The Strategist discussed accountability as a closely related concept, portraying such as benefiting from human involvement. Again, differentiating Aftonbladet from other companies was employed for illustration. *“There is no human being behind Facebook. [...] [At Aftonbladet] there is a human who has made a selection for you about what is going on in the world today and is commenting on it and prioritizing it for you. There is no one we can talk to at Facebook, it is basically just a computer doing it [...] At Aftonbladet, we have ‘the human touch’. There is someone behind [the brand]. You can go in and chat with us, they know that there are people sitting behind it and working”* (the Strategist).

Discussion

Researchers and practitioners in computational journalism put much hope to computational approaches providing relief to the pressured journalistic workforce. Our findings resonate with such ideas and suggest that in order for the promise of “everyday relief” through machine communication to be realized in journalism, a hierarchy between humans and machines should be manifested in the type of communicative roles they each are given. The results of the SEO-experiment provoked stakeholder reflections regarding what makes humans and machines distinct in their function as

communicators in the context of journalism, indicating that humans have the ability to explore the unknown (“the future”) while machines would maintain the status quo. We consider this to challenge the traditional anthropocentric view of communication by going beyond the *feasibility* of machines taking on communicative roles to distinguish what the functionality of those roles should be. Pairing these findings with recent research on computational journalism more broadly (e.g., Diakopoulos 2019b), we suggest a differentiation between *rule-based* and *knowledge-based* communicative roles in the newsroom. Such an approach could help decision makers in the newsroom practically evaluate potential use cases based on the degree to which they require abilities such as the generation of novel ideas or contextual understanding to be considered successful. Our findings indicate that such knowledge-based communicative roles are perceived as most suitably held by humans, suggesting that communicative AI employed for rule-based tasks in journalism would be met with comparatively more organizational optimism *and* be more suitable from a technical perspective.

Syvertsen et al. (2019) note how Nordic news organizations tend to use a “*public value vocabulary*” that reaffirms their civic responsibility in order to imply that the value of journalistic content extends beyond products. Our findings provide additional evidence of this. Related to the idea of news organizations’ self-assumed responsibilities, our research suggest that the presence of human communicators is perceived as a positive differentiation mechanism for journalistic organizations where human-curated and/or generated communication was suggested to build trust among readers. We consider how this perception may have limited the stakeholders’ appetite for continued exploration in the SEO-case, but also Aftonbladet’s readiness to leverage communicative AI more broadly. Building upon these findings, we suggest that in addition to assessing the rule/knowledge-based nature of the communicative task at hand, decision makers involved in evaluating potential use cases for communicative AI in journalism should consider whether the fact that there is a human performing the task at hand is or could be a value proposition in itself.

Both of the above discussed ideas (decision makers in journalism evaluating use cases based on the nature of communication, and the value of a humans performing it) will, as the SEO-experiment itself would have, benefit from the involvement of diverse competencies. In this regard, news organizations’ ability to make their collective competencies greater than the sum of their parts will be essential.

Gulliksen et al. (2020) suggest that mobility (meaning “*the ability to rapidly acquire knowledge, track change and acquire new knowledge, in conjunction with a motivation to contribute to digital development of society*”) is a core skill possessed by digitally excellent professionals across industries. The findings of this research suggest that such an ability is plausible to assume as desired but not yet possessed by the interviewed staff at Aftonbladet. And they would not be alone. Royal et al. (2020) argue that isolation between product and editorial departments in news media organizations is heightening tensions between them and suggest, as a remedy, a horizontal approach where departments work for mutual understanding of constructs and challenges. Similarly, Cornia et al. (2020) identifies an emerging “*norm of integration*” where commercial and editorial teams see the need for collaboration, adaptability,

and business thinking to effectively adapt to an increasingly challenging environment. We believe this study to add to an emerging body of research within digital journalism studies highlighting the need for “skills mobility” and increased collaboration, and more specifically, the challenges of getting such to work when engaging with advanced technologies as NLG.

Translating a word from one language into another requires one type of machine knowledge. Making a machine grasp its semantic meaning and generate new language (such as a SEO headline in a comparatively small language such as Swedish) is a different – and significantly more complex – task. There are undeniable technical challenges associated with the type of communicative AI implementation attempted through the experiment presented in this study, but we will leave discussion about them for other arenas.

Conclusions, Limitations and Future Research

This study assessed an experiment at the Swedish news publisher Aftonbladet using the machine learning technique of Natural Language Generation to create Search Engine Optimized headlines. While not achieving its goal, the experiment helps us understand some of the challenges associated with the ongoing transformation of journalism and how professionals in the field reflect on machines taking on new, communicative roles.

Conclusions

The findings of this study help us prepare for future communicative AI implementations in journalism. Three key ideas can be concluded in this regard. *First*, a functional distinction between rule-based and knowledge-based communicative tasks may aid decision makers in journalism to evaluate potential use cases for communicative AI. *Second*, investigations into the perceived value of humans performing a communicative task could help decision makers understand attitudes in their audience and organization. *Third*, there is a pressing need to foster interdisciplinary collaboration around communicative AI, where the skills of stakeholders involved in efforts employing the emerging technologies are shared with and understood by collaborators. As exemplified through this experiment, a lack of such skills mobility can limit organizations’ ability to develop and leverage computational approaches in the newsroom.

Limitations

This study provides rare insight into the process of exploring AI-tools in a news organization. In order to provide more general results, the experiment would have had to include a broader empirical base. Regarding the question of the feasibility of using NLG for generating SEO-headlines, increasing the number of articles tested in the technical experiment could have strengthened the conclusions. Similarly, the quality of the generated SEO-titles could have been evaluated by more than one person in order to avoid personal preferences or biases. However, as the remaining empirical data

(and the analysis of it) was not tightly connected to the results of the experiment but rather offered broader reflections on it, we consider the experiment to offer valid empirical basis despite these limitations.

Future Research

Our findings add to a growing body of illustrations of the challenges associated with using communicative AI in a legacy media organization (see for example, Jones and Jones (2019a)'s description of the BBC's use of chat bots). Our research contributes with a close-up look at a commercial media organization, and we encourage other researchers to study additional implementations in newsrooms of varying types, sizes and regions. Furthermore, the conclusions of this study would benefit from empirical testing and more in-depth research into the ideas of what constitutes a rule vs knowledgebased communicative task.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

ORCID

Agnes Stenbom  <http://orcid.org/0000-0002-0093-6859>
 Mattias Wiggberg  <http://orcid.org/0000-0001-9636-1545>
 Tobias Norlund  <http://orcid.org/0000-0002-4598-8652>

References

- Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. "Language models are unsupervised multitask learners." *OpenAI blog* 1 (8): 9.
- Allahyari, Mehdi, Seyedamin Pouriyeh, Mehdi Assefi, Saeid Safaei, Elizabeth D. Trippe, Juan B. Gutierrez, and Krys Kochut. 2017. "Text summarization techniques: a brief survey." arXiv preprint arXiv:1707.02268.
- Anderson, Christopher W. 2013. "Towards a Sociology of Computational and Algorithmic Journalism." *New Media & Society* 15 (7): 1005–1021. ISSN 14614448.
- Bahdanau, Dzmitry, Kyung Hyun Cho, and Yoshua Bengio. 2015. "Neural Machine Translation by Jointly Learning to Align and Translate." In 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings. <http://arxiv.org/abs/1409.0473>.
- Becker, Brian E., and Mark A. Huselid. 2006. "Strategic Human Resources Management: Where Do we Go from Here?" *Journal of Management* 32 (6): 898–925. ISSN 01492063.
- Beckett, Charlie. 2019. *New powers, new responsibilities: A global survey of journalism and artificial intelligence*. Polis, London School of Economics and Political Science. <https://blogs.lse.ac.uk/polis/2019/11/18/new-powers-new-responsibilities>.
- Bilton, Ricardo. 2018. Reuters' new automation tool wants to help reporters spot the hidden stories in their data (but won't take their jobs). Nieman Lab. <https://www.niemanlab.org/2018/03/reuters-new-automation-tool-wants-to-help-reporters-spot-the-hidden-stories-in-their-data-but-wont-take-their-jobs/>
- Broussard, Meredith, Nicholas Diakopoulos, Andrea L. Guzman, Rediet Abebe, Michel Dupagne, and Ching Hua Chuan. 2019. "Artificial Intelligence and Journalism." *Journalism & Mass Communication Quarterly* 96 (3): 673–695. ISSN 2161430X.

- Broussard, Meredith. 2018. *Artificial Unintelligence How Computers Misunderstand the World*. Cambridge, MA, USA: MIT Press. ISBN 9780262038003.
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan et al. 2020. "Language models are few-shot learners." arXiv preprint arXiv: 2005.14165.
- Carlson, Matt. 2015. "The Robotic Reporter: Automated Journalism and the Redefinition of Labor, Compositional Forms, and Journalistic Authority." *Digital Journalism* 3 (3): 416–431. ISSN 2167082X. <http://dx.doi.org/10.1080/21670811.2014.976412>.
- Coddington, Mark. 2015. "Clarifying Journalism's Quantitative Turn: A Typology for Evaluating Data Journalism, Computational Journalism, and Computer-Assisted Reporting." *Digital Journalism* 3 (3): 331–348. ISSN 2167082X.
- Cohen, Sarah, James T Hamilton, and Fred Turner. 2011. "Computational Journalism." *Communications of the ACM* 54 (10): 66–71. doi:<http://dx.doi.org/10.1145/2001269.2001288>.
- Cornia, Alession, Annika Sehl, and Rasmus Kleis Nielsen. 2020. "We No Longer Live in a Time of Separation': A Comparative Analysis of How Editorial and Commercial Integration Became a Norm." *Journalism* 21 (2): 172–190. ISSN 17413001.
- De Maeyer, Juliette. 2016. "Adopting a 'Material Sensibility' in Journalism Studies." In *The SAGE Handbook of Digital Journalism*, Chapter 31, edited by Tamara Witschge, C.W. Anderson, David Domingo, and Alfred Hermida. London: SAGE.
- Deuze, Mark, and Tamara Witschge. 2018. "Beyond Journalism: Theorizing the Transformation of Journalism." *Journalism* 19 (2): 165–181. ISSN 17413001.
- Diakopoulos, Nicholas. 2019b. *Automating the News: How Algorithms Are Rewriting the Media*. Cambridge, MA: Harvard University Press. ISBN 9780674976986.
- Diakopoulos, Nicholas. 2019a. "Towards a Design Orientation on Algorithms and Automation in News Production." *Digital Journalism* 7 (8): 1180–1184. ISSN 2167082X.
- Dick, Murray. 2011. "Search Engine Optimisation in UK News Production." *Journalism Practice* 5 (4): 462–477. ISSN 17512794.
- Dörr, Konstantin Nicholas. 2016. "Mapping the Field of Algorithmic Journalism." *Digital Journalism* 4 (6): 700–722. ISSN 2167082X. <http://dx.doi.org/10.1080/21670811.2015.1096748>.
- Duffy, Andrew, and Peng Hwa Ang. 2019. "Digital Journalism: Defined, Refined, or Re-defined." *Digital Journalism* 7 (3): 378–385. ISSN 2167082X.
- Eldridge, Scott A., Kristy Hess, Edson C. Tandoc, and Oscar Westlund. 2019. "Navigating the Scholarly Terrain: Introducing the Digital Journalism Studies Compass." *Digital Journalism* 7 (3): 386–403. ISSN 2167082X.
- Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114: 254–280. ISSN 00401625. <http://dx.doi.org/10.1016/j.techfore.2016.08.019>.
- Giomelakis, Dimitrios, Christina Karypidou, and Andreas Veglis. 2019. "SEO Inside Newsrooms: Reports from the Field." *Future Internet* 11 (12): 261. ISSN 19995903.
- Graefe, Andreas. 2016. Guide to Automated Journalism - Columbia Journalism Review.
- Grubenmann, Stephanie. 2016. "Action Research: Collaborative Research for the Improvement of Digital Journalism Practice." *Digital Journalism* 4 (1): 160–176. ISSN 2167082X. <http://dx.doi.org/10.1080/21670811.2015.1093274>.
- Gulliksen, Jan, Åsa Cajander, and Mattias Wiggberg. 2020. Digital spetskompetens—den nya renässansmänniskan: Genomlysning, definition, prognosverktyg och rekommendationer för framtida utveckling. Tillväxtverket https://tillvaxtverket.se/download/18.78563d971729d8c3af147501/1592477832340/DigitalSpetskompetens_Definition_Gulliksenetal.pdf
- Guzman, A.L. 2018. What is human-machine communication, anyway. Human-machine communication: Rethinking communication, technology, and ourselves, 1–28.
- Guzman, Andrea L., and Seth C. Lewis. 2020. "Artificial Intelligence and Communication: A Human–Machine Communication Research Agenda." *New Media & Society* 22 (1): 70–86. ISSN 14617315.
- Hollis, Martin. 2002. "Introduction: Problems of Structure and Action", in *The philosophy of social science: An introduction*. revised edition, 1–22. Cambridge University Press.

- Jones, Bronwyn, and Rhianne Jones. 2019a. "Public Service Chatbots: Automating Conversation with BBC News." *Digital Journalism* 7 (8): 1032–1053. ISSN 2167082X.
- Jones, Rhianne, and Bronwyn Jones. 2019b. "Atomising the News: The (in)Flexibility of Structured Journalism." *Digital Journalism* 7 (8): 1157–1179. ISSN 2167082X.
- Küng, Lucy. 2015. *Innovators in Digital News*. London: I.B. Tauris.
- Küng, Lucy. 2013. "Innovation, Technology and Organisational Change. Legacy Media's Big Challenges. An Introduction." In *A Multidisciplinary Study of Change What is Media Innovation?*, edited by Tanja Storsul and Arne H Krumsvik, Chapter 1, 9–14. Gothenburg: Nordicom, University of Gothenburg, ISBN 9789186523657. www.nordicom.gu.se.
- Latzer, Michael, Katharina Hollnbuchner, Natascha Just, and Florian Saurwein. 2016. "The economics of algorithmic selection on the Internet." In *Handbook on the Economics of the Internet*. Edward Elgar Publishing.
- Lewis, Seth C., and Nikki Usher. 2016. "Trading zones, boundary objects, and the pursuit of news innovation: A case study of journalists and programmers." *Convergence* 22, 5: 543–560.
- Lewis, Seth C., and Oscar Westlund. 2015. "Actors, Actants, Audiences, and Activities in Cross-Media News Work: A Matrix and a Research Agenda." *Digital Journalism* 3 (1): 19–37. ISSN 2167082X.
- Lewis, Seth C., Andrea L. Guzman, and Thomas R. Schmidt. 2019. "Automation, Journalism, and Human–Machine Communication: Rethinking Roles and Relationships of Humans and Machines in News." *Digital Journalism* 7 (4): 409–427. ISSN 2167082X.
- Skogen Lund, Kristin , and Ragnar Kårhus. Q1 2020 Results. Schibsted, 2020. <https://static.schibsted.com/wp-content/uploads/2020/05/06065305/Q1-2020Presentation.pdf>
- McNiff, J. 2016. *You and your action research project*. London: Routledge .
- Nedelkoska, Ljubica, and Glenda Quintini. 2018. *Automation, Skills Use and Training*. <https://www.oecd-ilibrary.org/employment/automation-skills-use-and-training/2e2f4eea-en>.
- Newman, Nic. 2011. "Mainstream Media and the Distribution of News in the Age of Social Discovery." *Reuters Institute for the Study of Journalism* 58.
- Organisation for Economic Co-operation and Development(OECD). 2021. Sweden - OECD Data, <https://data.oecd.org/sweden.htm>.
- Posetti, Julie. 2018. Time to step away from the 'bright, shiny things'? Towards a sustainable model of journalism innovation in an era of perpetual change. RISJ Reports. Reuters Institute for the Study of Journalism, Department of Politics and International Relations, University of Oxford.
- Raffel, Colin, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Li Wei, and Peter J. Liu. 2019. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. <http://arxiv.org/abs/1910.10683>.
- Royal, Cindy, Amanda Bright, Kirstin Pellizzaro, Valerie Belair-Gagnon, Avery E. Holton, Subramaniam Vincent, Don Heider, Anita Zielina, and Damon Kiesow. 2020. "Product Management in Journalism and Academia." *Journalism & Mass Communication Quarterly* 97 (3): 597–616. ISSN 2161430X.
- See, Abigail, Peter J. Liu, and Christopher D. Manning. 2017. "Get to the Point: Summarization with Pointer-Generator Networks." In *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, vol 1, 1073–1083. ISBN 9781945626753. <http://arxiv.org/abs/1704.04368>.
- Karlsson, Michael and Helle Sjøvaag, eds. 2018. *Rethinking Research Methods for Digital Journalism Studies*. London: Routledge.
- Steensen, Steen, and Laura Ahva. 2015. "Theories of Journalism in a Digital Age: An Exploration and Introduction." *Journalism Practice* 9 (1): 1–18. ISSN 17512794. <http://dx.doi.org/10.1080/17512786.2014.928454>.
- Syvrtsen, Trine, Karen Donders, Gunn Enli, and Tim Raats. 2019. "Media Disruption and the Public Interest." *Nordic Journal of Media Studies* 1 (1): 11–28.
- Tandoc, Edson C., Lim Jia Yao, and Shangyuan Wu. 2020. "Man vs. Machine? The Impact of Algorithm Authorship on News Credibility." *Digital Journalism* 8 (4): 548–562. ISSN 2167082X. [10.1080/21670811.2020.1762102](https://doi.org/10.1080/21670811.2020.1762102).
- Thurman, Neil, Konstantin Dörr, and Jessica Kunert. 2017. "When Reporters Get Hands-on with Robo-Writing: Professionals Consider Automated Journalism's Capabilities and Consequences."

Digital Journalism 5 (10): 1240–1259. ISSN 2167082X. <http://dx.doi.org/10.1080/21670811.2017.1289819>.

- Torbert, William R., Taylor, Steven S., 2008. Action inquiry: Interweaving multiple qualities of attention for timely action., In .London: SAGE, , 2: 239- 251,
- Torbert, William R. 2004. Action inquiry: The secret of timely and transforming leadership. San Francisco: Berrett-Koehler Publishers.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, L. ukasz Kaiser, and Illia Polosukhin. 2017. "Attention is All You Need." *Advances in Neural Information Processing Systems* 2017, 5999–6009. <http://arxiv.org/abs/1706.03762>.
- Vollstedt, Maike, and Sebastian Rezat. 2019. *An Introduction to Grounded Theory with a Special Focus on Axial Coding and the Coding Paradigm*, 81–100. Cham: Springer.
- Wagemans, Andrea, and Tamara Witschge. 2019. "Examining Innovation as Process: Action Research in Journalism Studies." *Convergence: The International Journal of Research into New Media Technologies* 25 (2): 209–224. ISSN 17487382.
- Wu, Shangyuan, Edson C. Tandoc, and Charles T. Salmon. 2019. "A Field Analysis of Journalism in the Automation Age: Understanding Journalistic Transformations and Struggles through Structure and Agency." *Digital Journalism* 7 (4): 428–446. ISSN 2167082X.

Appendices

A. Appendix 1

A.1. Model and training process

In the following paragraphs, we briefly introduce to the model architecture evaluated and reflected upon in this study. The rationale for publishing this information is to provide a basis for the stakeholder reflections presented in this article.

A.1.1. Models and Training Data

As an SEO headline should have some degree of similarity to the article it concerns, the task at hand in the experiment includes summarizing text. A model architecture that has proven particularly successful for that is called *the Pointer Generator* (See et al. 2017), which introduces a copy mechanism that effectively enables the machine learning model to learn when to copy words from input text (in this case, the journalistic articles). Furthermore, using the method of *transfer learning* through models such as the Transformer (Vaswani et al. 2017) is common in state-of-the-art NLG. Transfer learning refers to the process of first *pre-training* (parts of) a model on a different but similar task where more data is typically available, and then *fine-tuning* it on a dataset specific for the desired task. At Aftonbladet/Schibsted, more data were available on articles containing generic headlines and we could thus utilize transfer learning to first teach a model to learn how to generate such, and then fine-tune the model on a smaller dataset also including SEO-headlines (Table 3).

A.1.2. Step-wise Process

Training the model was a step-wise process (illustrated in Figure 1). First, the model was intended to learn how to write generic article headlines. To that end, we trained a Transformer model on the headlines of the *Schibsted-articles* dataset. In step 2, we conditioned the headline generation on the available article metadata by passing such (in the form of tags) to the Transformer. We "transfer learned" by re-using the model from step one and continued training

Table 3. Overview of training data.

Dataset name	Total dataset	Data from	Data content
Schibsted-articles	2,967,497 articles	Aftonbladet, Svenska Dagbladet and Omni	headline, body and metadata
Aftonbladet-seo-articles	83,317 articles	Aftonbladet	headline, body, metadata, SEO headline

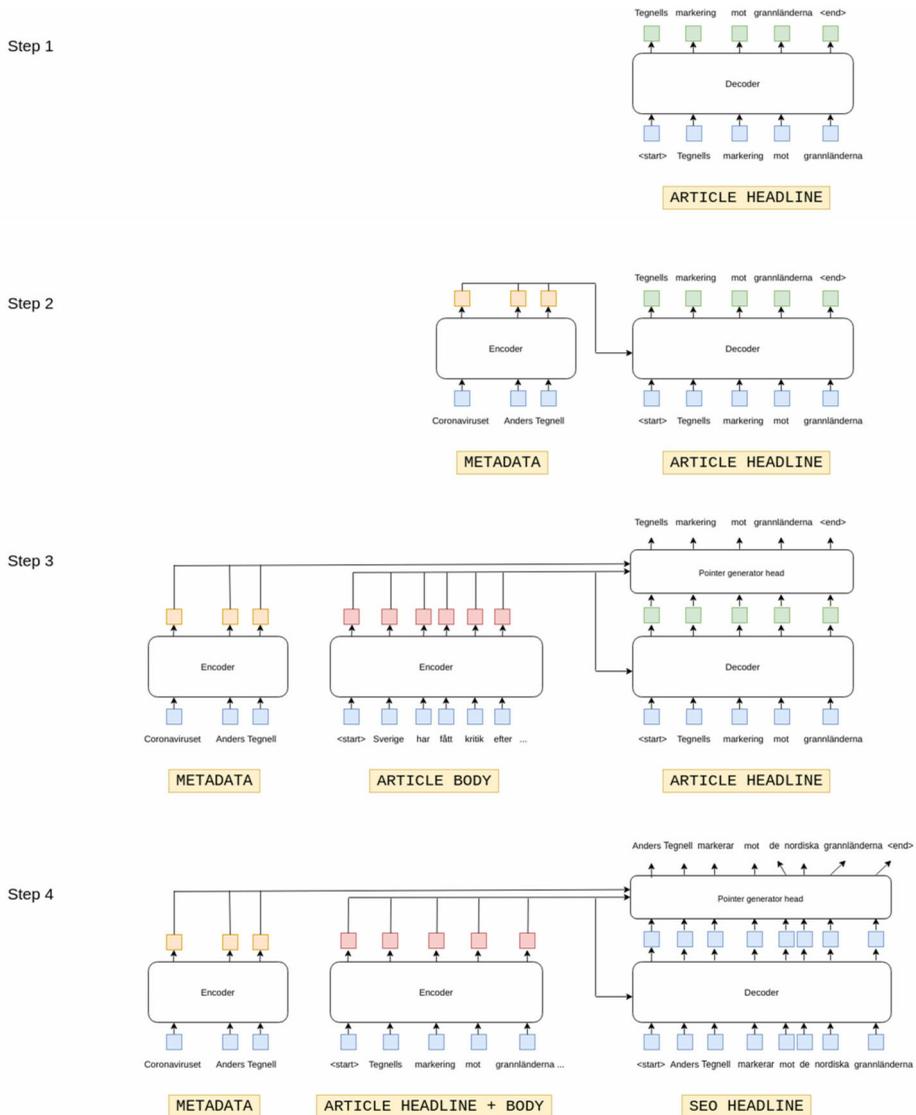


Figure 1. Visualization of the model architectures for each of the four training steps.

the Transformer model on the *Schibsted-articles* dataset, only now the headline generation was controlled by providing one or more article tags. In step 3, we added the article body as additional input to condition the headline generation further. The “copy mechanism” of the Pointer Generator was implemented, and we continued to “transfer learn” from step two and the *Schibstedarticles* dataset. The result was a model that could generate a generic headline for an article of which it was provided an article body and metadata. Finally, to achieve the goal of generating SEO-headlines, we trained this model also on the *Aftonbladet-seoarticles*. We provide the generic article headline as additional input by concatenating it with the body, and passed the resulting string as input.