Information Research - Vol. 30 No. iConf (2025)

Does Google dream of electric memes? From human to computational culture

Alexander O. Smith, Una Joh, and Jeff Hemsley

DOI: https://doi.org/10.47989/ir30iConf47350

Abstract

Introduction. We provide a cultural heritage informatic study of '*internet memes*': clusters of documents carrying a collective cultural meaning.

Method. Google's computer vision data from three memes are used to discuss the cultural heritage Google provides for the indexed web using Gini coefficients and word clouds.

Results. Our networks suggest that the more reducible sense-data is to a smaller number of labels, the less the internet needs to '*remember*' to see a '*meme*'. However, when labels are less strongly connected, human inspection is required to interpret if Google captures qualities of memes.

Conclusion. We show that Google can interpret aspects of memetic cultural heritage. However, it fails at understanding some cultural information about memes relative to predefined expectation. Computational tools like Google Cloud Vision can augment the heterogeneity of cultural heritage work. However closer inspection from experts is needed to interpret when computational methods represent cultural memory.

Introduction

Academic conceptions of '*internet meme*' typically assume a collection of files (e.g. jpeg or png files) that enact cultural memory: a defining interest of cultural heritage research (Assmann, 2008; Taylor, 1995; Viejo-Rose, 2015). As such, memes function as data of *cultural heritage informatics*: media imbued with the power of memory through iterative practices (Guan, 2008; Modrow & Youngman, 2023). Further still, memetics – the study of memes – studies dissemination, curation, classification, and organisation of culture (Dawkins, 1976; Shifman, 2014). In many cases cultural heritage informatics could easily borrow the name '*memetics*' without altering much. Yet, this sort of immediate association of research seems unlikely. Why?

The overlap seems conceptually intuitive, but underlying the intuition is a tension in disciplinary goals. Internet memetics regularly seeks to capture descriptions of '*new things*' in culture (Lankshear & Knobel, 2019; Smith et al., 2024): *disjunction*. Meanwhile cultural heritage seeks to find proper ways to canonise culture (Assmann, 2008) in relation to existing culture: *conjunction*. Internet memetics – despite conceptually being about memory and artefacts – is empirically resistant to the canonization of cultural connections by highlighting differences of interpretation (Lankshear & Knobel, 2019). Meanwhile, '*cultural heritage*' implies something inherited culturally. Cultural heritage pragmatically seeks to make '*something*' canon. Canonization of internet memes requires thoughtful responses within cultural heritage and memory institutions (Bratt & Smith, 2022; Rees, 2021; Tulloch, 2023), information studies broadly (Smith et al., 2024), but also within memetics (Pettis, 2021; Rogers & Giorgi, 2023).

In the internet research context, Pettis, Rogers, and Giorgi (2021; 2023) wonder how platform technology mediate, contextualise, and disjoin cultural memory. Instead, we turn to internet technology that conjoins the internet. Thus, we ask, '*How can Google search algorithms mediate, contextualise, and conjoin memetic memory*?' This question allows us to return to how to canonise the dynamics and heterogeneity of the collective memory of material culture. Thus, we redeem a fundamentally connected dimension of culture that memetic research seems to resist (Lankshear & Knobel, 2019) and the cultural dynamics of memory which *disciplinarily* defines heritage of cultural memory (Assmann, 2008).

To illustrate the issue, consider a brief motivating example. An '*internet meme*' is not a singular image-file nor a singular subjective memory. Rather, it is some arrangement (e.g., collection) of files, the way it moves through media, and the way it is remembered for every being who comes into contact with one or more of those image-files. Despite people frequently having claimed to have '*seen/made/shared a meme*,' this reduces a meme to a subjective interpretation of singular, isolated files. Likewise, saying '*I have curated a meme*' renders the meme dynamics invisible in the way they are technically and cognitively existing. Either case renders a meme without the context of how it is a meme, failing the conceptual objective of memes.

To properly '*curate a meme*' requires representing a meme inclusive of dynamic cultural meanings, dimensions, and implications. Researchers and cultural heritage practitioners must inseparably and accurately represent both (1) dynamic cultural memory and (2) an accurate description of its digitally mediated context.

This paper approaches the dynamics of memory for static image-files and the technologies that interpret '*memory*' for us as a kind of cultural heritage. Especially for memetics online, understanding how internet technologies interpret memetic content will shape the way we humans interpret culture within iterative expressions that emanate from memory practices. One of the most powerful technologies for the human interpretable internet is web search. Google, in particular, remains the most recognized search engine online (Bianchi, 2024). Understanding how Google *conjoins* the image-files that make '*a meme*' empirically closes a gap in connecting memes' static files to cultural heritage. In particular, the data that establishes regularities in web search and information retrieval are systematically organized through iterative, aggregative, networked, and collective habits of search. While perhaps one might suggest that '*humans*' are absent in studying Google search, it is data humans 'invisibly' produce through search that provides a significant portion of Google's regularized data for information retrieval.

Literature review

Memetics is usually considered a kind of datafication of cultural information (Smith et al., 2024). Richard Dawkins (1976) coined the term meme to describe a unitary bit of cultural information which inherited from '*mind-to-mind*.' He exemplifies this through architectural '*arches*' as a cultural artefact in the world which is fundamentally connected to the mental representation of an '*arch*'. Later, Daniel Dennett integrated this into his theory of mind (1991/2017). Dennett illustrated how cognitive connections pass these idealised bits of culture by way of information, using a myth about how Marco Polo brought pasta from China to Italy through memory and documentation.

These early accounts of '*meme*' were critiqued and re-theorised for digital contexts by Limor Shifman (2014) and others (see, Smith et al., 2024). Yet, memes do actively entangle artefacts and memory (Smith & Loewen-Colón, 2024), but few researchers have empirically shown how '*internet memes*' and memory connect.

Often '*mental*' perspectives – cognition, phenomenology, memory, knowledge, and so forth – are buried just under the data of technical objects (Rogers & Giorgi, 2023) or communication data (Shifman, 2014). Yet, meme's information conjoins minds, albeit with more nuanced and multi-mediated ways than Dawkins' or Dennett's reductive '*mind-to-mind*' descriptions. As contemporary memetic research clarifies, there are often a plurality of information technologies between '*mind-to-mind*' mappings.

To redeem mental processes of memetics, how can we begin to understand what '*memory*' is when the data are outside of human minds, and instead are in the world as artefacts? Previous research has suggested a relationship between cultural heritage and memes (Demsky, 2021; Modrow & Youngman, 2023). However, empirical interests of memetic cultural heritage (Demsky, 2021; Petrassi, 2024; Rees, 2021; Tulloch, 2023) usually investigate human-centric experiences without the nuance of technical mediations that conjoin minds. As such, there is a gap between internet memetic technological arguments and cultural heritage's research on individual human experiences. Complex digitally mediated ecologies of memes (Shifman, 2014) and cultural heritage need to be re-joined to explain memetic memory (Smith & Loewen-Colón, 2024).

When memetic data are collections of image-files, the cultural heritage task is clarified. Relatedly, memory activates through these files. Despite being corollary, digital files and mental processes are not reducible to each other. Shifman (2014) suggests '*social mindsets*' entangle inseparably with these files but remains open to how this occurs. She defines memes as '*content*, *form*, *and/or stance*' which must travel across the internet by way of many users. Content is the '*idea*,' form is the meme's '*media*', and stance is how the information enables unique contextual meanings for individuals as the meme is communicated. Despite the intentional entanglement of these concepts, internet memetics have tended towards turning '*content*' of a meme into a plurality of meanings (Lankshear & Knobel, 2019; Smith et al., 2024) which tempts researchers into reducing Shifman's definition to merely digital media's immediate context (Wiggins, 2019). Doing so renders the transmission and dynamics of cultural heritage characterised within Shifman's definition inaccessible. To redeem memory, Smith, and Loewen-Colón (2024) create initial steps in re-theorising memory for both Shifman and Dawkins.

While lots of theorisation of transmission of memetic information exists (Cappella et al., 2015; Dawkins, 1976; Hull, 2000; Smith & Hemsley, 2022; Smith & Loewen-Colón, 2024), observing transmission in digital contexts is a difficult task without near exact copies of artefacts such as '*retweets*' or '*share*' buttons (Sperber, 2000). This is partly why '*viral events*' were distinguished from '*memes*' in digital contexts (Nahon & Hemsley, 2013; Shifman, 2014). However, information transmission and meme might have nuanced connections if we could conceptualise what information constitutes a meme in some statistically regular image-file data. To fill the empirical gap between materiality of image-files and the '*content*' of memetic memory, we must start by seeing information selection across image-files.

Theoretically, the relation of these image-files and memory have been recently updated as being connected to '*traces*' (Smith & Loewen-Colón, 2024). Information selection of digital and cybernetic technologies are often called *trace data* (Geiger & Ribes, 2011; Østerlund et al., 2020). Trace data signifies movement or *selection* (Cappella et al., 2015; Heylighen, 1993) of information. Studying memetic traces empirically remains incompletely described, yet previous literature has plenty of suggestions.

Prior work suggests interpreting memetic data traces are digitally mediated and technologically situated (Rogers & Giorgi, 2023; Smith et al., 2022). However, Smith et al. suggests trace data can be understood through the datafication of Google image web search technologies (Smith et al., 2022; Smith & Hemsley, 2022) and across digital technologies more generally (Smith & Loewen-Colón, 2024). Google both provides traces of web indexing, but also artefactual traces '*within*' each artefact. We distinguish *artefact* from jpeg or png files: memetic *artefacts* are data commonly identified across files. When Google identifies a frog or a cat in two different image-files, that frog or cat operates as *artefactual trace data* which Google considers culturally relevant for web search, indexing, and so forth. Google '*traces*' a cat to a bunch of other cat artefacts. If Google considers this data relevant for web search, then it is this data that informs our cultural experience online.

By using Google as a method of recovering dynamics across image-files by tracing artefacts across them, we ask the following questions:

> RQ1. What do Google's digital traces suggest about the conjoining of memetic image-files as technologically mediated cultural heritage? How are different meme's image-files similar and/or different?
>
> RQ2. Do Google's digital artefact relations seem correlated to the digital traces provided by a memetic collection of image-files on their own through our initial interpretation?

RQ1 considers technologically mediated conjunction using computational network methods and compares across three qualitatively selected memes based on what the authors deemed structurally different kinds of traces. RQ2 considers these three structurally different meme image-file collections to interpret similarity and differences (Smith & Hemsley, 2022) of artefactual traces across the files. The first question considers what the technology conjoins as the regular semiotic '*structure*' of each collection. The second asks what we expected would be conjoined in relation to what actually was conjoined. Together these questions suggest how technologically mediated conjunction might be different from expected conjunction. We will explain how we provided these differences across three memes in the next section.

Methodology

With interest in what data traces that sustain collective/distributed internet memory, we use relevant technologies that provide this datafication following Smith et al. (2022) who use a cyber-semiotic interpretation (Cannizzaro, 2016) of memetic traces. As such, our research design makes use of Google data gathered from Google Cloud Vision's API. This provides us data that informs Google about what image files are affiliated with text search and reverse image search. These data play a large role in popular web search, shaping our digital cultural environment.

However, we are also interested in how a '*meme*' that people curate into a collection of image documents '*belong together*'. For this we leverage the cultural '*authority*' of Know Your Meme's (KYM's) image collections as other researchers have done (Katz & Shifman, 2017; Smith et al., 2022; Smith & Hemsley, 2022). KYM curates image-files into thousands of '*meme entries*,' each of which operate as cultural heritage collections for '*a meme*.' The following section shows how we use KYM's meme entries to gather cultural heritage data from Google.

Data collection

Before getting data from Google, we downloaded three meme entries: Doge, Hello There, and Loss. Doge typically contains a Shiba Inu dog in a variety of contexts. '*Hello There*' is a meme based on character dialogue in Star Wars between Obi-wan Kenobi saying, '*Hello There*' and General Grievous responding, '*General Kenobi!*' However, as can be seen below they are somewhat visually correlated, but they are not as ubiquitous in their visual data as Doge. Loss is a pattern '| || || |_' in some literal or translated way. Example image-files of these meme entries can be seen in respective order in Figure 1 below:

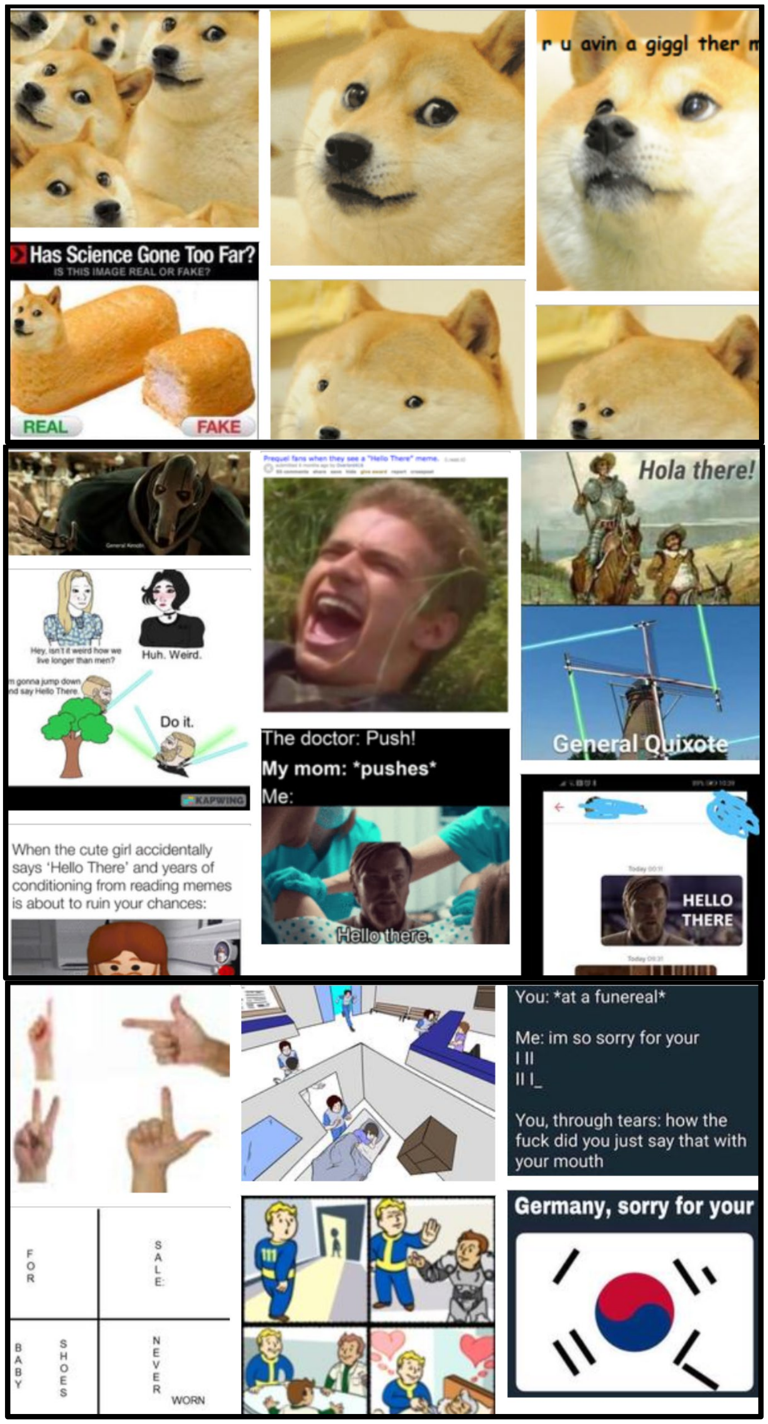

Figure 1. Examples of Doge, Hello There, and loss screen captured from KYM

These were chosen because of visual differences and the way they are culturally remembered through traces that conjoin the image-files together. Through visual inspection, the authors noted the collection of image-files appear to belong together. '*Doge*' is a collection of files which almost exclusively belong to the collection because they have the iconic face of a photoshopped Shiba Inu dog that is mostly used as a reaction image. It is often used to signify an emotion like an emoticon. '*Hello There*' image-files each contain some subset of a larger set of regular co-occurring artefacts within each file. Finally, Loss is unique in that it contains a geometrically abstract pattern. While all the images contain this pattern, the pattern itself seems to require abstract imagination to observe it. While it might be obvious that Google can observe a Shiba Inu and assume an association to '*Doge*' in text search, assuming an abstract pattern like Loss would be akin to searching for '*Hello There*' in Google and expecting that the text alone will return lots of Star Wars results.

While there are perhaps many '*different kinds*' of memes than this, we could not find a complete taxonomy of memes that tries to interpret them through inspections of trace data, although it seems methodologies which might enable this exist (Smith & Hemsley, 2022). Although there are highly context dependent taxonomies of things like '*kinds of memetic nonsense*' (Katz & Shifman, 2017) which are not within the scope of the present study's goal.

Thus, after having collected the three meme entries, we used Google Cloud Vision API to return the relevant image-file features. Our relevant *features* were '*labels*' and '*web entities*'. *Labels* are tags that Google's cloud vision attaches to images. For example, if Google sees a Doge affiliated image-file, it will likely return the label '*dog*'. Labels are more immediately observed visual artefacts within the image file and serve more as immediate sense data. *Web entities* are search terms or Google's '*inferences*' that are mostly likely associated with the images in web search. For example, if Google recognizes a dog in a Doge image, then likely it will return web entities such as 'Shiba Inu', '*Doge*', etc that are more indicative to cultural heritage, and thus the connective '*memory*' of the image file. Labels proxy as '*sense-data*' while web entities proxy for technologically sustained '*memory*': i.e., the cultural heritage sustained by Google.

Data analysis protocol

After having collected the image-files from the three meme entries and the Google Cloud Vision data, we have a set of human readable data that reflects what visual artefacts are in each image document, but also aggregately within each meme collection. To interpret what artefacts, get '*selected*' within each meme collection, we interpret the labels and the web entities as '*nodes*'. However, instead of including all the detected features as nodes, we choose features that have more than a 0.5 confidence score, which shows how certain Google Cloud Vision is about the detection of that feature. From these nodes, we construct two sets of networks of co-occurring features: one for labels, and another for web entities for each meme entry.

In feature-to-feature networks, the weight of an edge represents the number of common images-files where two labels or web entities appeared together. This is useful because, low co-occurrences, such as those that only occur once (edge weight = 1) can be weighted as less important to a meme entry than those that occur within 10 images (edge weight = 10) for example. Weighted graphs enable us to calculate the '*strength*' of nodes, a metric that is the sum of the weights of the edges connected to a node (Barrat et al., 2004). This serves as an alternative measure of centrality of feature data to a meme entry, accounting not only for the number of connections, as in degree centrality, but also for the intensity or importance of those connections. The strength of the features enables us to '*quantify*' the visual similarity and differences between two images and help us describe the 'fuzziness' of semiotic references as we understand them in internet memes.

From the strength value of each node, we calculate the Gini coefficient. The Gini coefficient is a measure of inequality in a distribution, ranging between 0 and 1, where 0 corresponds to perfect equality. In our networks, the Gini coefficient would be close to 0 if all the features have the same strength, and close to 1 when there is a large disparity among the strengths. Even though each meme entry has a different number of nodes (the number of features that appear in the images of each meme entry), we can still compare the Gini coefficients because it is a relative measure of inequality that is independent of the network's size.

Results

There are two types of networks: (1) feature networks and (2) image networks. Feature networks represent how the features (either label or web entities) are connected to each other based on common images that include both entities. Conversely, image networks show how individual images-files are connected based on common entities that are detected in both images.

Feature networks

We calculated the strength of all the nodes in the network, which was used to produce a Gini coefficient by using a '*Lorenz curve*' (Hu & Wang, 2005; Kirillov & Panov, 2022). The Lorenz curve provides the basis for calculating the Gini coefficient by using the ratio of the area between the Lorenz curve and the line of perfect equality to the total area under the line of perfect equality. The resulting Gini coefficient allows us to compare inequality in the distribution of strength values. That is, a lower Gini coefficient suggests that strength values are quite evenly distributed across features. Conversely, a higher Gini coefficient shows that a few nodes (features) have very high strength values—meaning they share a lot of commonalities with others—while most nodes have lower strength values. Then the Lorenz Curve was produced by ordering the ratio of nodes by their strengths, seen in Figure 2 for each meme.

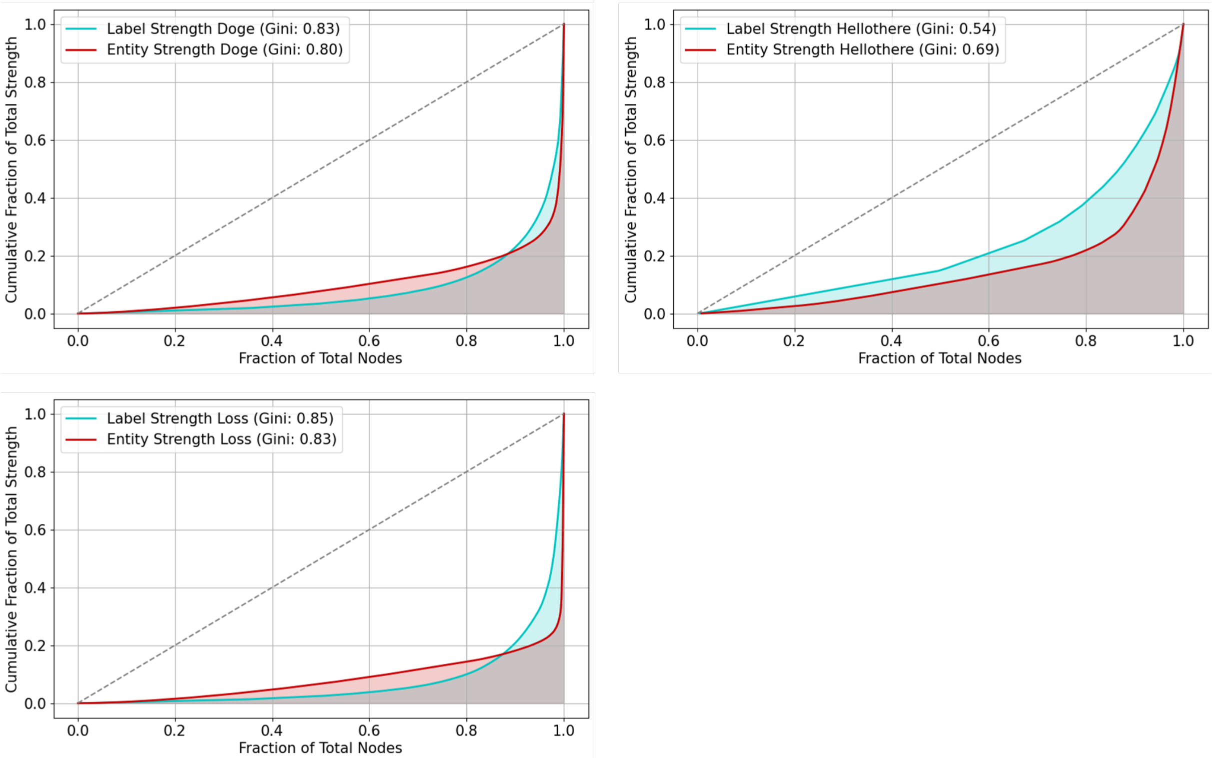

Figure 2. Lorenz curve and Gini coefficient of node strength: comparison between labels and entities

What Figure 2 shows us is a comparison of each meme's label and entity features' Gini coefficient. As the figure shows, Doge and Loss are relatively more hierarchical in their labels than their entities. That is to say when Google interprets the labels, it extends those labels to entity data that are less hierarchical. This means Google '*sees*' the image-files as having more specific visual features, but it has less '*memory*' of what those image-files are about specifically. The opposite is true of Hello There. Google 'sees' the image-files of Hello There, and 'remembers' more specifically what it is about.

Answering RQ1, we find that Google captures the cultural heritage of Hello There in higher specificity than Loss or Doge because of how it extends to more specific cultural contexts with the strength of particular entity labels' meanings to the meme. This can be seen in a word cloud for Hello There. We see that from the more dominant labels, there are a lot of references which may or may not obviously relate to the cultural '*meaning*' of Hello There as a meme. However, as we move to the entities, the cultural context of Hello There becomes more clear in that Google sees that it is related to Star Wars and its characters for example.

However, Doge and Loss, from Figure 2 appear remarkably similar. This is counterintuitive based on our visual inspection. In order to distinguish if Google actually sees and '*remembers*' them differently we look at what it sees relative to what Google remembers by considering the entities names themselves. We can also extend what we just saw from Hello There and see what happens to the most '*strong*' features as we move from Doge's and Loss' labels to entities in Figure 3, below.

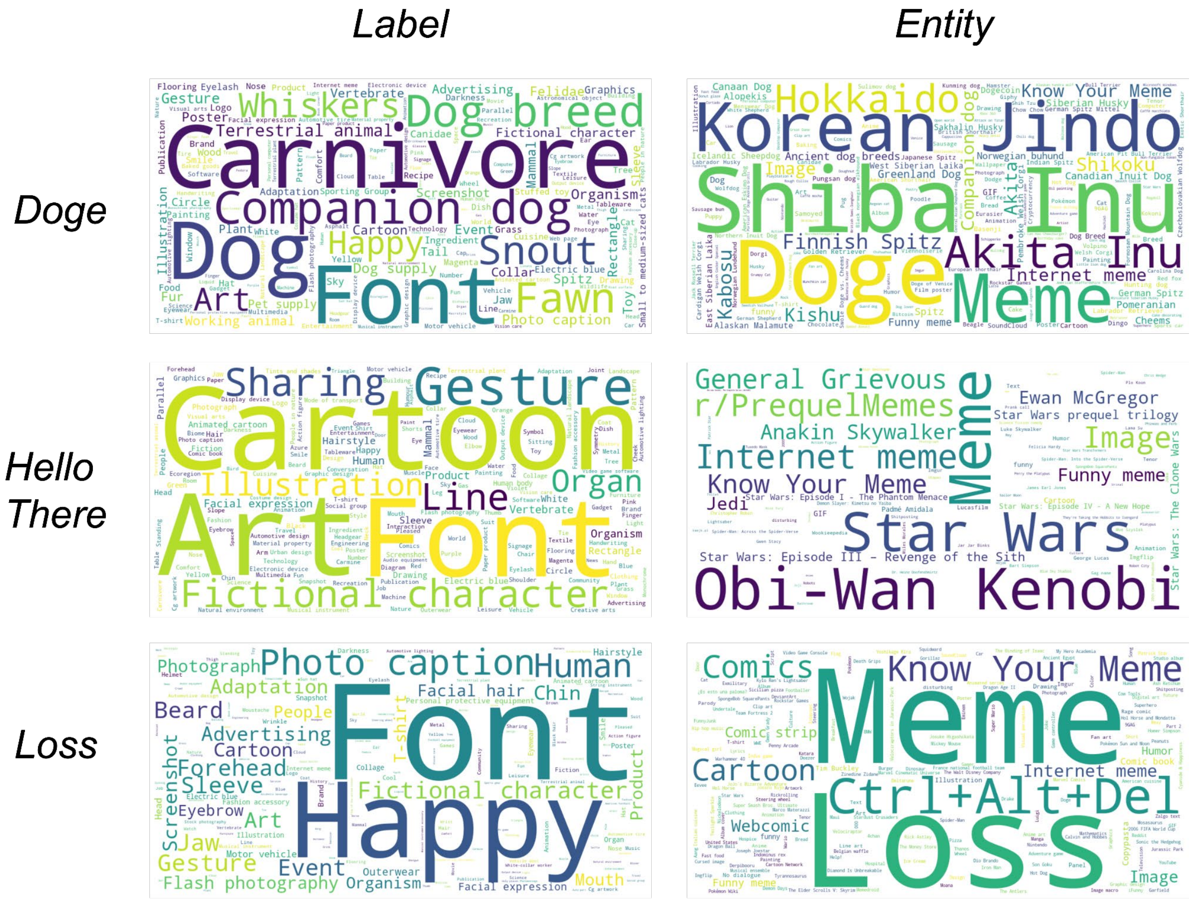

Figure 3. Word cloud by strength of feature

This figure suggests that for Doge, the entities it remembers are much closer to what we expected. Google '*sees*' text and a dog, and '*remembers*' it is a 'meme' of a '*shiba inu*' called '*Doge*'. That is, while perhaps Google knows more of what '*Doge*' is informationally about based on cultural expectations. It remembers that it is seeing a Doge meme and many different aspects core to Doge, which is the core of the meme. At first glance, it appears as though the same is true for Loss. However, recall what makes Loss culturally relevant is more about a pattern than what we see here. While Google '*remembers*' it is a meme called 'Loss' it captures more information about its source material: Ctrl+Alt+Del, a cartoon or webcomic. Or alternatively it particularly generically refers to Know Your Meme in a nonspecific way. If it captured what makes Loss, the meme entry, '*belong together*,' one of the primary words would be something like '*pattern*' or some text representation of Loss: '| || || |_'. Google trace data's lack of mapping to human memory of what makes Loss culturally specific is what was expected. It is this way that Doge and Loss are different from each other in Google's sustained trace data of these memes. Google's data suggests it knows what is culturally relevant about Hello There and Doge in Figure 3. However, Figure 3 shows us that the critical bit of information that makes Loss different in the kind of meme it is from the other two is not captured. Google merely reiterates generic data about the meme. This answers RQ2.

Discussion

Hello There

Our results do suggest strong reasons to think that Hello There is significantly different from Doge or Loss. While Doge and Loss are consistently extremely hierarchical, Hello There is consistently the opposite. In terms of Google's features, this is because of the combinatorial arrangements of what allows something to fit into '*Hello There*.' As such, there is more cultural content to describe Hello There. As we see in the Figure 3, despite there being more cultural content, it seems to summarise what we expected. Hello There comes from a narrative interpretation provided from Star Wars that has regularly been translated into other languages, narrative contexts, and iconographic aspects of Star Wars. That is, Google 'sees' things that are highly correlated to Star Wars, and thus Google connects the correlates as a regularised '*digital memory*'. In our authorial minds, it is intuitive that the meme entry would have observable connections to the narrative. As our Gini coefficient suggests, it is that when the labels are organised together, somehow Google algorithmically predicts relations to web entities based on digital regularities. Google expects those labels to extend towards terminology that is associated with Star Wars, lightsabers, Obi-Wan Kenobi, General Grievous, and so forth. So, the web entities focus on aspects of the Star Wars story in its slightly more hierarchical interpretation of the meme. Since we '*know*' the story of Star Wars, this is sensibly intuitive because the data we typically remember are digitally representable. However, the actual computational method of doing this is inaccessible to us, but it must be computational and algorithmic, and thus digital. So, while it requires a large number of references for a digital cultural actor to understand Hello There as a human cultural phenomenon. However, the word cloud shows that it does not remember the '*narrative*' call-response, '*Hello There*', followed by '*General Kenobi!*' which is what the meme ultimately is about. So, while it focuses on the surrounding narrative, it does not seem to recognize it fully.

Doge

Doge's word cloud also maps well to our memory of it based on our interpretation of what is core to Doge: The Shiba Inu and surrounding characteristics. As a meme that's about bodily gesturing and facial characteristics, it captures something very similar to that. However, we noticed the Gini coefficient suggests Google '*remembers less*' about the meme. This is perhaps because Doge has '*less*' information about it other than the particular dog in the image. Doge culturally floats into and out of many different cultural contexts. People post it everywhere with no particular '*linguistic*' goal, but rather it has an expressive goal like how someone might use an emoticon. This is contrary to '*Hello There*' which is a call-response sort of joke from Star Wars. That is, we know Google knows Doge because the Doge meme culturally is a dog and not much else: a floating signifier. '*Hello There*' is known through the nodes that strongly hold it together: the combinatorial features. However, the word cloud suggests that it does not understand the '*call-response*' cultural aspect of the meme.

Loss

Loss, as expected, is particularly difficult to capture by Google both by Gini coefficient and by what we culturally expected to be in the word cloud. This is because what makes '*Loss*' work is abstract, and difficult to reduce to language. Simply networking the image-files together creates a less strong, cohesive idea of what the meme centres around. And, as such, it remembers the source material and where the meme is mostly documented more so than what the meme means or is about. Google merely repeats its context back to us without content.

Conclusion

In total, we see that Google's '*memory*' of memes as collections of image-files in some cases performs better at capturing what we expected of memetic cultural heritage, and sometimes it does not. In particular, Google understands what '*Hello There*' and '*Doge*' are about, however we have to see these two memes in different ways. Hello There, is understood more as a network of features by Google where Doge only seems to be understood through the dog, but what holds the meme together otherwise is not known to Google. Loss, however, seems to lose context in both cases. Essentially Google thinks '*Loss*' is a webcomic or a generic meme on Know Your Meme called Loss. It does not give us any indication of its cultural relevance.

What our study suggests is that certain parts of cultural memory are not reducible to computationally reducible visual data. More context is needed to understand what makes a meme important. We see this in each of our meme's network data. However, we do see that Google can provide a partial understanding of what artefactual trace data frequently ties them together. As such, we can begin to see what information selection is occurring across collections. However, in the case of Hello There and Doge, it seems Google cannot point to exactly what gets selected. That is, Google's attempt to describe Doge leads it to various descriptions of Doge and not the cultural contexts in which Doge is used. It is even less precise about Loss. On the other hand, Hello There is more precise in its network, but it is less precise about its narrative content other than it is '*funny*'.

In relation to memetics, we have seen that some memes are understood through computational artefact networks through Google vision features. However, we have also given some insights into where it is limited. For example, it does not understand when cultural context is not within the data, or when the cultural context is strongly implied through probabilistic references in the case of Loss and Hello There.

Our literature review suggests that Google can actually play a significant role in what gets sustained in cultural heritage, and thus can shape cultural memory. As such a significant part of our cultural knowledge, how it understands or misunderstands this culture plays a part in a computationally sustained cultural heritage as trace data. Memes, when understood as '*particulates*' of documented culture, provide a way to be granular about understanding computationally driven tools like Google Cloud Vision's usefulness or limitations in sustaining the relevant parts of cultural memory. In order to be clear about what Google generates for cultural heritage, connections between cultural expectations of those that '*know*' and an understanding of how the qualitative aspects of cultural heritage are bridging with computational tactics of making sense of them. While our results suggest that Google is '*good*' at understanding certain parts of memetic meaning, we find that it needs to be augmented through human interpretation to make sense of a lot of what it provides. Google, in particular, can significantly aid in bridging new cultural heritage together however, since it pulls together a lot of information in the heterogeneity of the indexed web, so understanding such tools helps augment limited subjective views of cultural heritage work. Future work is needed to understand what computational aspects can be fully leveraged in augmenting cultural heritage work and what should be delegated to humans.

About the authors

Alexander O. Smith is a PhD candidate of Information Science and Technology at Syracuse University. Their research is focused on memetics: the study of memes. Alexander is forwarding a novel theoretical framework for memetics which synthesizes older memetic theories with internet meme studies. Affiliated work can be found at *Convergence*, *Journal of Documentation*, ASIS&T, and the IEEE big data conference computational archive science workshop. More expansive research interests on (cross-)platform dynamics, network cultures, and classification can be found at *Policy and Internet*, *Social Media + Society*, eCSCW, and AoIR.

Una Joh is a Ph.D. student at the School of Information Studies at Syracuse University, specializing in computational social science. With a background in economics from her bachelor's and master's studies at SKKU in Korea, she conducts research on the flow of information and the dynamics of public opinion on social media, taking into account its interaction with news media. Her research aims to extract, track, and predict information flows and opinion dynamics on networks that model interconnected media environments. To achieve this, she utilizes a variety of methods, including natural language processing, network analysis, and agent-based modeling.

Jeff Hemsley is the Interim Dean at the School of Information Studies at Syracuse University. He earned his Ph.D. from the University of Washington's information school. His research is about understanding information diffusion, particularly in the context of politics, in social media. He is co-author of the book Going Viral (Polity Press, 2013 and winner of ASIS&T Best Science Books of 2014 Information award and selected by Choice magazine as an Outstanding Academic Title for 2014), which explains what virality is, how it works technologically and socially, and draws out the implications of this process for social change.

References

Assmann, A. (2008). Canon and Archive. In A. Erll & A. Nünning (Eds.), Cultural Memory Studies: An International and Interdisciplinary Handbook (pp. 97–108). De Gruyter. https://doi.org/10.1515/9783110207262.2.97

Barrat, A., Barthélemy, M., Pastor-Satorras, R., & Vespignani, A. (2004). The architecture of complex weighted networks. Proceedings of the National Academy of Sciences of the United States of America, 101(11), 3747–3752. https://doi.org/10.1073/pnas.0400087101

Bianchi, T. (2024). Online search market worldwide—Statistics & Facts (Online Search). Statista. https://www.statista.com/topics/1710/search-engine-usage/

Bratt, S., & Smith, A. O. (2022). Evolutionary Archives: The Unlikely Comparison of GenBank and Know Your Meme. 2022 IEEE International Conference on Big Data (Big Data), 2490–2499. https://doi.org/10.1109/BigData55660.2022.10020929

Cannizzaro, S. (2016). Internet memes as internet signs: A semiotic view of digital culture. Σημε$ιota$ømega$τκ$\acute\eta$-Sign Systems Studies, 44(4), 562–586. https://doi.org/10.12697/SSS.2016.44.4.05

Cappella, J. N., Kim, H. S., & Albarracin, D. (2015). Selection and Transmission Processes for Information in the Emerging Media Environment: Psychological Motives and Message Characteristics. Media Psychology, 18(3), 396–424. https://doi.org/10/gf8x87

Dawkins, R. (1976). The Selfish Gene. Oxford University Press.

Demsky, J. (2021). That Is Really Meme: Nazi Pepe the Frog and the Subversion of Anglo-American Holocaust Remembrance. In J. Demsky (Ed.), Nazi and Holocaust Representations in Anglo-American Popular Culture, 1945–2020: Irreverent Remembrance (pp. 105–125). Springer International Publishing. https://doi.org/10.1007/978-3-030-79221-3_7

Dennett, D. C. (2017). Consciousness Explained. Little, Brown. (Original work published 1991)

Geiger, R. S., & Ribes, D. (2011). Trace Ethnography: Following Coordination through Documentary Practices. 1–10. https://doi.org/10/fqntnc

Guan, K. C. (2008). Oral Histories in the Making of Our Memories and Heritage. https://doi.org/10.1163/156853108X327100

Heylighen, F. (1993). Selection Criteria for the Evolution of Knowledge. Proc. 13th Int. Congress on Cybernetics (Association Internat. de Cybernétique, 524–528.

Hu, H.-B., & Wang, L. (2005). The gini coefficient's application to general complex networks. Advances in Complex Systems, 08(01), 159–167. https://doi.org/10.1142/S0219525905000385

Hull, D. L. (2000). Taking memetics seriously: Memetics will Be what we make it. In R. Aunger (Ed.), Darwinizing culture: The status of memetics as a science (pp. 43–67). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780192632449.003.0003

Katz, Y., & Shifman, L. (2017). Making sense? The structure and meanings of digital memetic nonsense. Information Communication and Society, 20(6), 825–842. Scopus. https://doi.org/10.1080/1369118X.2017.1291702

Kirillov, B., & Panov, M. (2022). Measuring internal inequality in capsule networks for supervised anomaly detection. Scientific Reports, 12(1), 13575. https://doi.org/10.1038/s41598-022-17734-7

Lankshear, C., & Knobel, M. (2019). Memes, Macros, Meaning, and Menace: Some Trends in Internet Memes. Communication and Media Studies.

Modrow, S., & Youngman, T. (2023). Theorizing Cultural Heritage Informatics as the Intersection of Heritage, Memory, and Information. Proceedings of the Association for Information Science and Technology, 60(1), 666–671. https://doi.org/10.1002/pra2.836

Nahon, K., & Hemsley, J. (2013). Going Viral (1 edition). Polity.

Østerlund, C., Crowston, K., & Jackson, C. (2020). Building an Apparatus: Refractive, Reflective, and Diffractive Readings of Trace Data. Journal of the Association for Information Systems, 21(1), 1–22. https://doi.org/10.17705/1jais.00590

Petrassi, D. (2024). Memetic Aesthetics between Artistic Production and Cultural Heritage. This is why we must consider memes the multimodal artifacts of the 21st century.. Southern Semiotic Review Issue, 19. https://doi.org/10.33234/SSR.19.3

Pettis, B. T. (2021). Know Your Meme and the homogenization of web history. Internet Histories, 0(0), 1–17. https://doi.org/10/gmkh47

Rees, A. J. (2021). Collecting Online Memetic Cultures: How tho. Museum and Society, 19(2), 199–219. https://doi.org/10/gmfdcw

Rogers, R., & Giorgi, G. (2023). What is a meme, technically speaking? Information, Communication & Society, 0(0), 1–19. https://doi.org/10.1080/1369118X.2023.2174790

Shifman, L. (2014). Memes in Digital Culture. The MIT Press.

Smith, A. O., & Hemsley, J. (2022). Memetics as Informational Difference: Offering an Information-Centric Conception of Memes. Journal of Documentation, 78(5), 1149–1163. https://doi.org/10.1108/JD-07-2021-0140

Smith, A. O., Hemsley, J., & Tacheva, Z. Y. (2024). Reinforming memes: A literature review of the status of memetic information. Journal of Documentation, 80(4), 1003–1021. https://doi.org/10.1108/JD-10-2023-0218

Smith, A. O., & Loewen-Colón, J. (2024). Memetic memory as vital conduits of troublemakers in digital culture. Convergence. https://doi.org/10.1177/13548565241262421

Smith, A. O., Tacheva, J., & Hemsley, J. (2022). Visual Semantics of Memes: (Re)Interpreting Memetic Content and Form for Information Studies. Proceedings of the Association for Information Science and Technology, 59, 800–802. https://doi.org/10.1002/pra2.731

Sperber, D. (2000). An Objection to the Memetic Approach to Culture. In R. Aunger (Ed.), Darwinizing Culture: The Status of Memetics as a Science (1st ed., pp. 163–173). Oxford University Press.

Taylor, H. A. (1995). “Heritage” Revisited: Documents as Artifacts in the Context of Museums and Material Culture. Archivaria, 8–19.

Tulloch, B. J. (2023). Memes to an End: Why Internet Memes Matter to Information Research. Libri, 73(4), 293–305. https://doi.org/10.1515/libri-2021-0098

Viejo-Rose, D. (2015). Cultural heritage and memory: Untangling the ties that bind. Culture & History Digital Journal, 4(2), e018. https://doi.org/10.3989/chdj.2015.018

Wiggins, B. E. (2019). The discursive power of memes in digital culture: Ideology, semiotics, and intertextuality. Routledge.