Behind Closed Doors: The Human Remains Trade within Private Facebook Groups

Shawn Graham, Katherine Davidson and Damien Huffer

The existence of a thriving trade in human remains online is facilitated by social media platforms. While much of this trade is conducted in fully public forums such as e-commerce platforms, the retail website of bricks-and-mortar stores, public personal and business pages on social media, etc., there also exist numerous private groups using the affordances of various social media platforms to buy, sell, and share photographs of human remains. This article describes a case study of four private Facebook groups featuring people who buy and sell human remains, to explore how the discourses of the trade may be different when not made in public. Using a close-reading approach on the text of posts and threaded conversations, and associated visual similarity analysis of the accompanying photographs, we observe, among other things, a strikingly 'more professional' approach, shibboleths and patterns of behaviour that serve to create group identities. We analyse posts made over a seven-week period across the selected private groups in the run-up to the 2023 holiday season. Given the issues of privacy raised by studying private groups, we also experiment with a locally hosted large language model to see if it could classify discourses meaningfully without the intervention of a researcher having to read the original posts. This case study might also serve as a model for other kinds of research investigating the reception of various archaeological topics that might be discussed and understood differently in private versus public venues.

1. Introduction

When people meet in private Facebook groups to buy or sell human remains, how do they talk about the remains they have or seek? What characterises the discourses, and are these discourses different from how people interested in this trade behave in public online forums?

Over the last several years, we have been studying the public posts made by individuals, participants, traders and enthusiasts relating to human remains across

multiple platforms like Instagram, TikTok, Facebook and Marktplaats (Davidson *et al.* 2021; Graham and Huffer 2020; Graham *et al.* 2020b; Graham *et al.* 2020a; Graham *et al.* 2022; Huffer and Graham 2017; 2018; 2023; Huffer *et al.* 2021; Huffer *et al.* 2019; Huffer *et al.* 2022). When posts are made on these open public platforms there can be no reasonable expectation of privacy, according to research ethics guidelines provided by the Canadian Social Science Humanities Research Council of Canada (Tri-Council 2022, 48). Nevertheless, we have taken measures to respect the privacy of people involved in this trade because the goal of this research is to understand why people buy and sell human remains, what physically owning remains does for these people, how sharing photographs of human remains shapes community aesthetics, goals and values and, to a lesser degree (insofar as it is possible to do so working solely from photographs), to attempt to understand what broad-level populations human remains bought and sold may derive from.

Having dealt with public-facing posts of this trade synoptically in our 2023 volume, *These Were People Once* (Huffer and Graham 2023), this article shifts to explore the human remains trade as evidenced in private Facebook groups (daily illicit human remains trade activity globally). These groups encourage participation in the trade through validating people's interest in buying and selling human remains as a hobby like any other hobby, while the recommendation algorithm brings buyers and sellers together in a supportive atmosphere. In Davidson *et al.* (2024) we discuss our protocol for obtaining and joining such groups and for studying the materials in a way that respects the privacy of participants and allows us to create aggregate data through which we can obtain a globalised view of what goes on inside these groups. (We refer to these groups as 'private' throughout but, given that the barrier to entry is usually not that onerous, a more accurate phrasing would be 'semi-private' as they pertain to Facebook. We could also note that 'privacy' online is a continuum that exists as a function of a user's technical facility intersecting with the affordances of a given platform.) Following that protocol, here we will not reproduce photographs from these groups, nor will we quote text verbatim from these posts. Once the data have been collected, we use a qualitative analysis approach to close-read the posts, where we read and re-read the posts and conversations over and over, creating, collapsing, and assigning categories until each post is described with appropriate tags describing the discourse.

As an experiment, we also used a locally hosted large language model (so that the data never had to leave the researcher's personal computer), prompted to categorise the discourses present in a post or series of related posts (a thread). We then compared these results to tags and emergent discourses identified through our close reading. The point of this experiment is to see how the language model matched the expert categorisation provided by the researcher, for if the locally hosted large language model could do a good job of categorisation (however measured), it would imply a way to study private posts *en masse* in a way that would completely obscure individual identities from the researchers, meeting the requirements for research ethics in our and many other universities, as well as funding bodies. Finally, we measure the photographs accompanying the posts for visual similarity (via image embeddings in a neural network model), as a way of assessing the degree of cross-posting between separate, private groups.

2. Method

2.1 Collecting

We described this study to, and sought permission from, the Carleton University Research Ethics Board Panel A in autumn 2023; permission was granted on 14 November. All data for this article were collected between 14 November and 31 December 2023. We describe in detail the process of obtaining that ethics clearance and the argument and rationale used in Davidson *et al.* ([2024](#)). We refer the reader to that article but, to summarise, the key element involved compartmentalising our research process so that researchers involved in analysing the posts could not inadvertently 'out' a person involved in these groups. DH collected and anonymised the posts with screenshots and screen blurs; SG automatically transcribed the posts using object character recognition (OCR) software and manual error checking; SG and KD categorised the posts through iterative re-reading and tagging with themes using a customised version of Ryan J.A. Murphy's ([2021](#)) 'Integrated Qualitative Analysis Environment' for the free Obsidian.md knowledge management software.[1]

For this article, a keyword search of Facebook's 'Groups' feature surfaced a surfeit of groups related to the human remains trade and the collection of 'oddities'. The goal was to mimic a casual new collector's first foray into this world, and so only groups in the first results provided by Facebook were explored. Four that seemed to capture the variability adequately (based on our own subjective experience of the wider trade) were selected for study. One was devoted explicitly to the buying, selling or trade in human remains only, while the other three were explicitly linked to the appreciation of 'oddities' in general. 'Oddities' is a standard keyword that will surface a variety of materials including human remains (Huffer and Graham [2017](#)). While there are thousands of posts and conversations, we focused on those posts or threads that undeniably were about buying, selling, or trading human remains.

DH led the data collection using an established 'lurker' account built for research purposes and selected the groups. All groups explored and considered for inclusion used Facebook's tools for semi-automating group membership moderation ('Admin Assist'). These tools might pose questions or require new members to check a box confirming that they will abide by codes of conduct, however the group sets them. The sheer number of members in these groups (often in the thousands and tens-of-thousands range) can be taken as an indication of how pro-forma this barrier to entry is. Nevertheless, if the answer to a question posed would have required us to lie, then we did not progress any further with a given group.

Once admitted to membership, post selection followed a 100% sampling strategy. This meant that every post that was undeniably offering human remains of any sort (including cremains, as well as art or jewellery incorporating human remains) for sale or trade was screen captured using a built-in snipping tool. As per our Carleton University Research Ethics Board approved protocol (Project #120176), the names of the four groups are known only to DH and have not otherwise been shared. All files

are stored on removable media within folders labelled generically: Group 1, Group 2, Group 3 and Group 4. Filenames incorporate metadata indicating the generic group label and the date when the post was made. Screenshots captured the entire thread of a discussion at the time of data capture. All screen shots were entirely anonymised via blacking/blurring out all possible identifiers such as names, profile pictures, store names, URLs, or living individual's faces appearing in photographs. Screenshots were then sent to SG who transcribed the text of captions and discussions automatically using object character recognition (OCR) software.[2]

2.2 Reading

The resulting transcriptions were then read using a grounded theory approach (Baumer *et al.* 2017; Charmaz 2017, 38; Dillon 2012; Moore *et al.* 2019; Nelson 2020). The second and third researcher in the process (SG and KD) would read each transcription to characterise the discourse(s) present with short descriptive tags or identify some kind of theme. As patterns emerged, they would iteratively cluster the characterisations, reading deeply and across posts to identify themes and reduce redundancy. Sometimes themes or tags would be subsumed within larger ones; sometimes it made sense to split a general theme into more descriptive sub-themes to capture the nuance. This close reading approach contrasts with the distant-reading approach we have used in previous studies using tools such as Topic Models or Word Embeddings (Huffer and Graham 2017; Graham and Huffer 2020). We then counted and visualised the tags/themes and their co-occurrences across posts creating a networked visualisation of tags that 'go together' to generalise conclusions about how human remains are discussed in these private groups.

2.3 An experiment in computerised reading

As an experiment (see below) we also used a locally hosted Large Language Model (LLM) Mistral 7B to 'read' each post/conversation and suggest its own categorisations of the discourse themes. LLMs excite interest because when connected to a chatbot interface, they seem to be able to generate sensible-looking text in response to their interlocutor's queries, creating an illusion of intelligence and agency. LLMs are statistical representations of words in the context of their use in written expression; they are a massive elaboration of the same basic processes that underpin familiar technologies like predictive text autocompletion. LLMs emerge from a variety of research agendas, one of which is machine translation where a model learns parallel representations of ideas across different languages. Here we are not interested in text-generation or 'chatting' about our data, but rather, we wish to use the model's representation of language to 'translate' the data we have into a more concise representation. We imagined that our qualitative assessment of the discourse surrounding each post was a translation task from the verbose natural language of the posts to a more concise analytic language of our tags. Because this was a locally hosted language model on a personal computer, no data left the investigator's machine.[3] We then explore the degree to which these auto-suggestions match our close reading. If the congruence was good, then this might imply a method that could

be used to scale up and study large amounts of textual data obtained from (semi-)private social media while respecting privacy concerns. We used Simon Willison's 'llm' python package for working with locally hosted Large Language Models (Willison [2023]). Since each conversation/post was contained within a single text file in our project folder, Willison's package allowed us to quickly write a bash script to feed each conversation/post through the LLM, appending the output to a csv file for examination.

2.4 Image analysis

For image analysis, we were interested in determining if specific categories of human remains were being posted to different groups. Visual similarity between images can be measured if the images are turned into vectors using an image embedding model and the distance between vectors measured. The images that accompany a post were transformed into an image embedding model using a locally hosted version of the Inception 3 model from Google (Google [2024]). Inception 3 is a neural network representation of thousands of images, where each image is represented as a multi-dimensional vector or list of numbers. A vector describes a direction in the space captured by the neural network. A 'vector' that might be familiar would be the conventional way of expressing the location of points in cartographic space as a list: [45.424721, -75.695000]. This is the 2-dimensional vector describing the location of Ottawa. An image vector contains hundreds of more dimensions. We encode the images we are studying into that space using the model, turning them into lists of numbers expressing distance along each dimension; images that are visually similar end up occupying the same direction in that space. Thus, we can measure the cosine distances between embeddings, which permits us to create a dendrogram of visual similarity to explore whether or how visual tropes span the different groups: do different groups have different visual grammars?

3. Results and Discussion

3.1 Qualitative assessment of the posts

Between 14 November 2023 (when our Ethics Clearance was issued) and 31 December, we observed 303 separate 'conversations' across four groups. By conversation, we mean an original post and any follow-up comments (creating a thread) made on that post.

- Facebook Group 1: 32 conversations
- Facebook Group 2: 126 conversations
- Facebook Group 3: 48 conversations
- Facebook Group 4: 97 conversations

We iterated over the groups several times applying tags, adding, removing, coalescing, and breaking up several 'themes' that seemed to capture the nature of the discourse, and the human remains being discussed. The tags are present in all the

groups, to different degrees, and in different combinations. Some tags might indicate the *how* of a conversation; others might represent the *what*, etc. Table 1 lists all tags across all groups and the number of times that particular tag was used. A single post might carry a number of different tags: we then understand a particular kind of discourse from the combination of tags present. For instance, 'PriceAsEmoji', 'NoShippingToCertainStates' and 'SearchAvoidance' suggests one kind of discourse where selling is framed as something clandestine, while 'HowToPurchase', 'SalesProcessUpdate', and 'PricesGiven' implies a certain professionalisation of salesmanship. Both of these might further combine with various kinds of stories concerning the origin of the remains, providing deeper nuance still. We represent those combinations of tags in the network graphs below in the discussion for each group.

Table 1: Themes identified across all posts collected, with frequency represented in the 'Count' column

3.1.1. Facebook Group 1

In this group, prices were very nearly always given, and there was usually an attempt to hide the price from (presumably) text-based searching or data-mining by using emojis to display the digits. Nevertheless, the OCR process we used recognised the digits without any issue. Another habit of this group we labelled 'searchAvoidance': the deliberate misspelling of key words, especially the word 'human'. In nearly half of the conversations, the initial post always contained the word 'human' rewritten to include spaces or sometimes the digits '00' in place of the letter 'u'. Pronouncing 'human' as if it was spelt with a double-o would place a stress on the vowels similar to the way the space-faring 'Ferengi' species were depicted to pronounce the word across multiple Star Trek television series. We do not think that this is accidental, but a kind of in-joke that also serves to dehumanise what is being offered for sale. The convention in this group seems to always note whether or not shipping is included in the price, and then to invoke almost a protective spell: 'no shipping to the verboten trio', or similar, always meaning Louisiana, Georgia, and Tennessee.

We can see four main patterns of discourse employed in this group by mapping the themes to the original posts as a kind of network, theme to conversation. Then, we can transform this network so that themes are connected to themes by virtue of the number of conversations; by measuring the patterning of connections, we can visualise groups of themes, as in Figure 1.
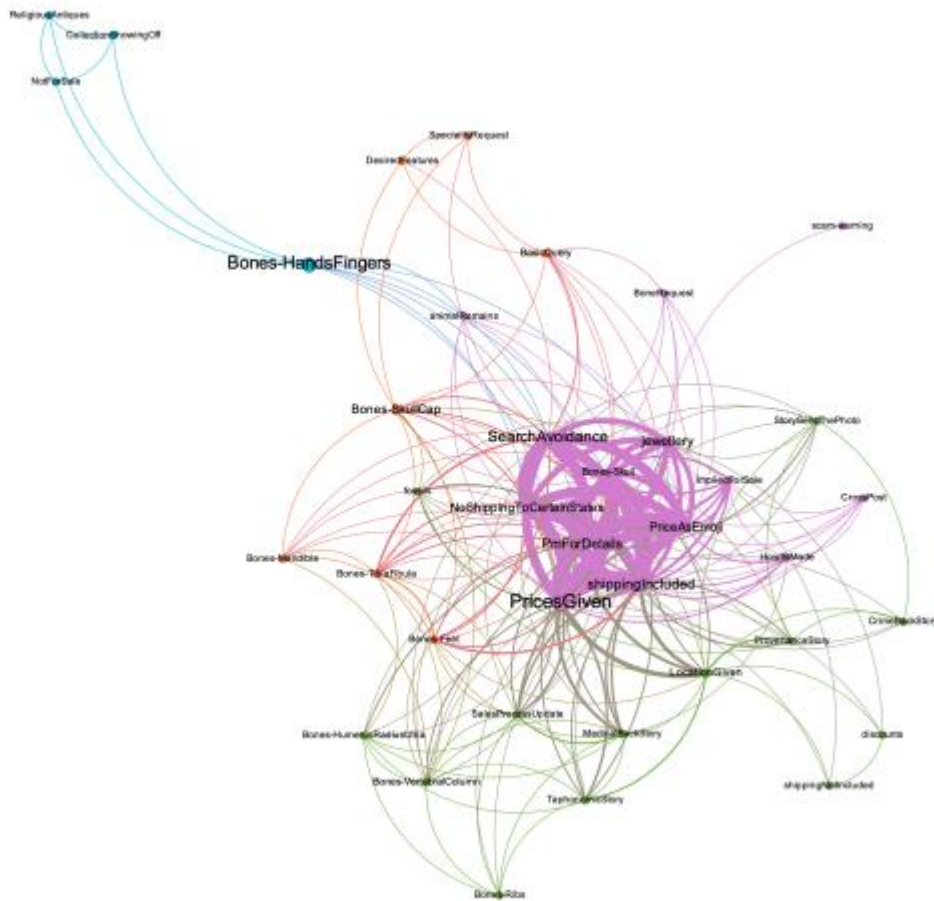
Figure 1: Facebook Group 1, tags connected to tags by virtue of appearing in the same conversations. The thicker the lines, the greater the number of co-occurrences. Tags that appear as most central are scaled accordingly larger. Colour indicates 'modules' or subgroups based on similarity of interlinkages

Facebook Group 1 has four 'modules' or subgroups within it. The largest subgroup concerns the *strategies* for how the participants achieve sales - they include prices via emoji, they include shipping, they misspell words deliberately, and they are often selling jewellery that incorporates human remains. Related posts in this broad discourse include discussions of how a particular piece is made, and sometimes the implication is that other things depicted are also possibly for sale, if a person were to reach out with a private message. Conversations here sometimes also discuss what constitutes a 'good' seller. There are very few 'meta' discussions (that is, discussions about the legality of trading, or how to make a trade, or the kinds of things that might be suitable for this group). There are a few requests for objects of desire, but the vast majority of conversations are like any classified ads marketplace: here is what I am selling, here is what I want, here is how to contact me.

The second broad pattern of discourse in this group circulates around *what* is being bought and sold, and the stories that support the sellers' contentions - stories about how they obtained the remains in the first place, or stories about the physical points-

of-interest or desirability of specific remains on offer (patina being a constant signifier). Backstories attesting to the former use of remains on offer as medical or dental specimens are frequently encountered. 'Shipping not included' is also a stock phrase within this group, used when sellers elaborate on the origins of the remains, perhaps to retain more of the sales value. Some animal remains or fossils are also sold alongside human materials.

Another constellation of discourses that interconnect here can be referred to as the 'speciality requests', where individuals who desire a particular kind of human remains make their initial posts. Others might chime in with what we would call 'basic queries' about how sales are made, or what kinds of remains might be on offer - the most desirable being human skulls.

A final constellation of discourses in Facebook Group 1 are the 'not for sale' posts, wherein individuals share pictures of their collections. These posts often centre around the display of religious antiques and sometimes quite ornate displays, for example, articulated hands.

Facebook Group 1 as a whole is structurally more diverse than the other groups, although the corpus sampled from it has the fewest conversations. By structurally diverse, we mean that the conversations tend to remain thematically distinct, whereas even though the other Facebook groups have more themes across more conversations, those themes tend to tie together into smaller clumps.

3.1.2 Facebook Group 2

This group had more conversations across the month and a half we observed and had a broader variety of individual themes than the other groups, but the way those conversations tie together thematically only suggests two broad discourses (Figure 2). It was a group more tightly focused on buying and selling.
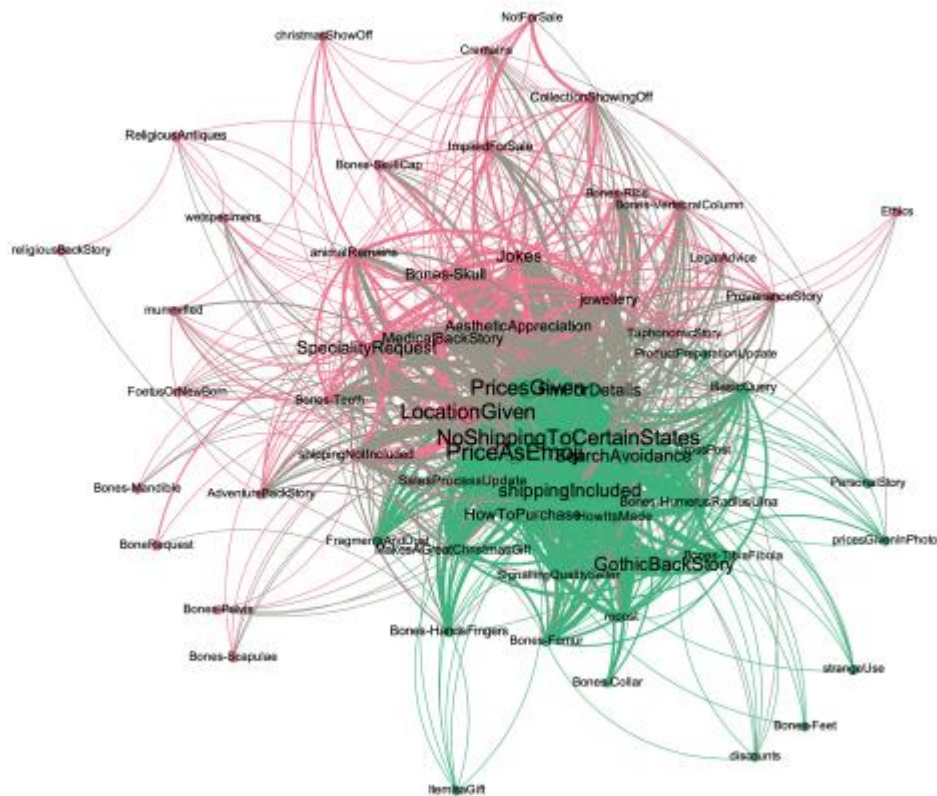
Figure 2: Facebook Group 2 tags connected to tags by virtue of appearing in the same conversations. The thicker the lines, the greater the number of co-occurrences. Tags that appear as most central are scaled accordingly larger. Colour indicates 'modules' or subgroups based on similarity of interlinkages

The first large pattern of discourse is very similar to what we saw with the largest discourse cluster for Facebook Group 1, but with more detail. There is a distinct 'How To Purchase' discourse, and a 'How It Is Made' discourse, in addition to the emojis and search avoidance that we saw before. Indeed, conversations that participate in these discourses seem to be signalling that they are a quality seller, a seller one can trust, a purveyor of fine materials because they go into detail not only about the origins of the materials, but also how they have been mounted or otherwise displayed. These individuals or organisations are engaging with social media in a way that could be considered 'textbook' social-media commerce best practices. They update posts or conversations when something has been sold. Sometimes they make a new post with an item that has been sold, marking it as such, in a way reminiscent of how real-estate agents will display past sales on their signage or websites, the implication being 'This is the kind and quality of wares that I can move'. The Christmas season was also a marketing tactic here, with variations of 'makes a great Christmas gift!' or 'show your love!' appearing closer to 25 December. (Between the

25th and the 31st, a related theme of showing off the human remains received as gifts also emerged.)

The other constellations of themes observed in this group included the *kinds* of items for sale and how those items were described in terms of the 'story that sells' (Huffer and Graham [2023]). Speciality requests tended to be for human skulls or wet specimens (which are understood by most participants as being scientific specimens preserved in formaldehyde, though not by all). There was some showing off of collections (explicitly marked as 'not for sale', though in several posts this was clearly a coy wink to the viewer, indicating that perhaps something might be sellable after all). A related theme was 'aesthetic appreciation', where other members of the group would coo over the depicted remains - and almost immediately, inappropriate jokes would also be told.

3.1.3 Facebook Group 3

The main cluster of discourses in this group again are very similar to the other two Facebook groups - patterns of key phrases to signal something for sale (and to protect oneself from observation), patterns of discourse to signal that a seller is rather higher quality than most, either in their conduct or their wares (Figure 3).

The other two main clusters of discourse for this group were grounded in showing off collections (and the aesthetic appreciation of the same, with the concomitant jokes), and discussions around the legality of the trade or how to actually make a sale. In this latter set of discourses, prolonged discussions of various laws in various jurisdictions sometimes took place. These conversations intersected with the 'I'm a quality seller' discourse, demonstrating knowledge of the law (see Dundler [2021] and Breda [2023] on 'personas' within the antiquities and human remains trades), but as conversations progressed the lack of certain professional legal knowledge within the group would cause the conversation to loop and circle fruitlessly.
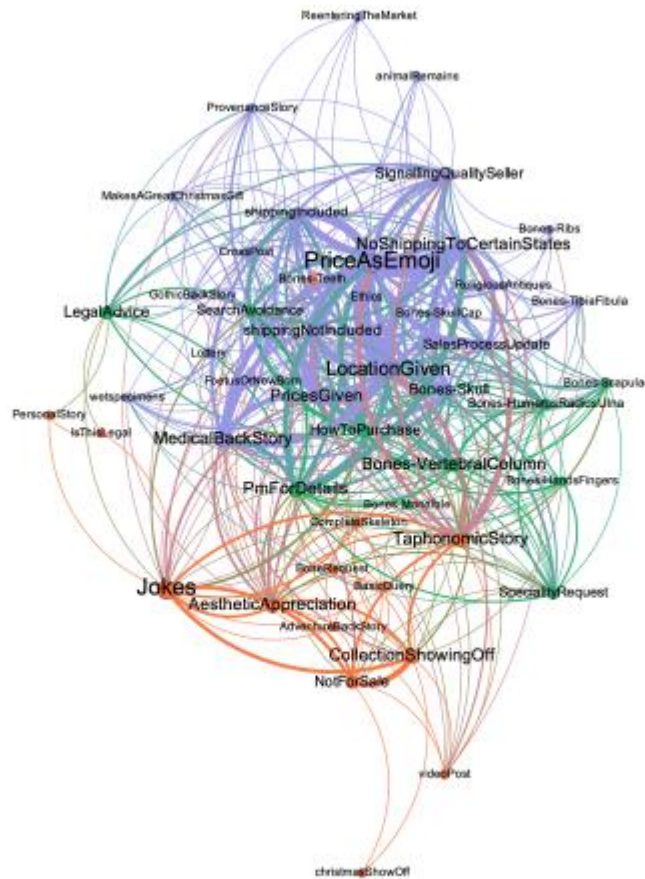
Figure 3: Facebook Group 3 tags connected to tags by virtue of appearing in the same conversations. The thicker the lines, the greater the number of co-occurrences. Tags that appear as most central are scaled accordingly larger. Colour indicates 'modules' or subgroups based on similarity of interlinkages

3.1.4 Facebook Group 4

This group had two main constellations of discourses, with the main being (as present in lesser degrees in the other groups) the shibboleths and key phrases that 'protect' the vendor as they make the sale (Figure 4). The main constellation of how-to-sell discourses, including the ways 'I'm a reputable dealer' could be signalled, here involved in-depth conversations of the merits of various digital tools to support refunds or how to arrange for insurance in case of damage during shipping. This group also tended to have a higher preponderance of discourses related to Christmas gift-giving (and the desirability of human remains or wet specimens for such things), as well as a greater degree of 'back stories' to signal the bona fides of the person selling or the desirability of the object. In this we see an element common across groups where a smaller subset of the participants *teaches* other participants how to act, how to 'consume' human remains. Another aspect of the discourses in this particular group were the cross-posts or re-posts from the other three groups we observed. That four groups chosen at random should have this much cross-posting

(making the same initial post in different groups, with the same picture and the same language) suggests that there is probably quite a large degree of interlinkages among human remains-specific or 'oddities' focused groups in general (far beyond the four surveyed here), which has implications for how 'the trade' as a whole learns to inhabit this space, as well as how established buyer-seller relationships can cross platforms.
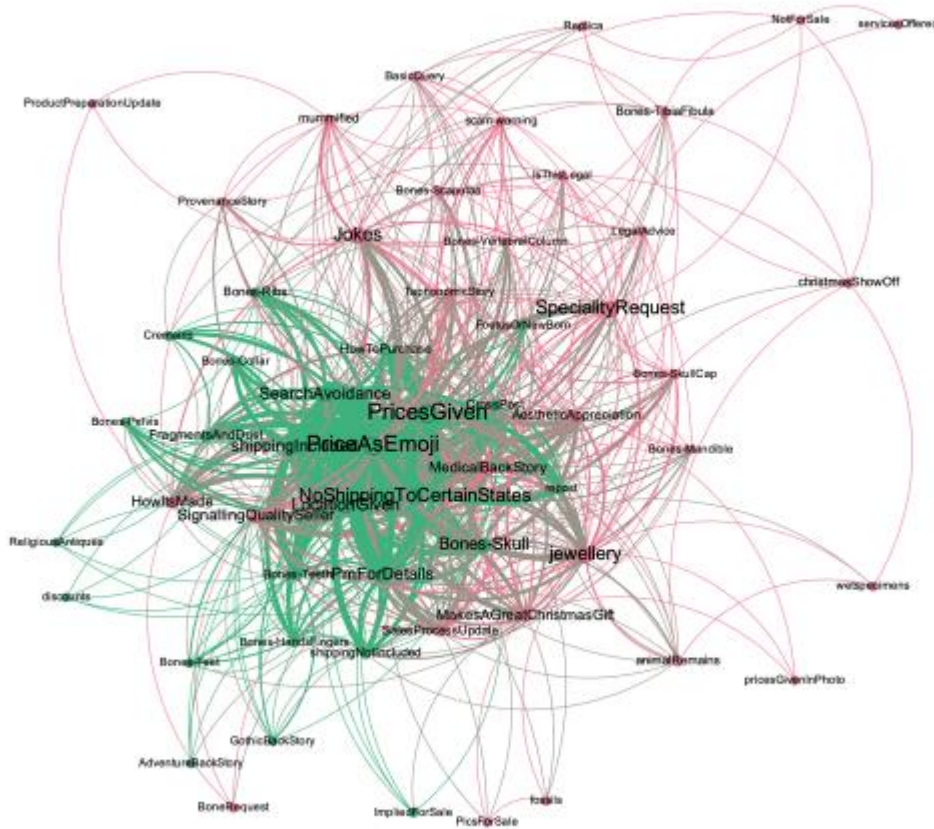

Figure 4: Facebook Group 4 tags connected to tags by virtue of appearing in the same conversations. The thicker the lines, the greater the number of co-occurrences. Tags that appear as most central are scaled accordingly larger. Colour indicates 'modules' or subgroups based on similarity of interlinkages

3.2 Large language model assessment

With a locally hosted LLM, there is a possibility of adding an additional level of anonymity to the source data by removing one level of human inspection. Note that we are not proposing using a LLM in the service of some kind of spurious 'objectivity' or from a desire to increase speed or 'efficiency'. We are proposing that in situations where one is dealing with sensitive information, a computerised, localised (not connected to the internet) approach *might* offer privacy protection. By getting the computer to do the reading and assign its own descriptive tags we might reduce the possibility of one of the investigators revealing private information (and indeed

perhaps protect the investigator from the cumulative impact of reading such posts). Such an approach would respect the privacy of users that are engaged in an activity that exists in a moral and legal grey zone (Huffer and Graham 2023, see chapter 4).

This experiment was suggested to us via the success of a project unrelated to the bone trade, where one of us used an LLM's ability to understand the structure of written English to identify statements of fact about the antiquities trade and then arrange them as subject, verb, object triples (Graham *et al*. 2023; that task was similarly an act of translation or distillation). Through careful exploration and consideration of what others have reported across a variety of venues (for the most part, on personal blogs or social media platforms), we found a prompt that seemed to guide the model well on the desired task while not prefiguring the result. Using Simon Willison's 'llm' package for working with local models systematically, we ran the following prompt using a consumer grade Mac Mini M1 with 16 gb of ram:

```
$cat transcribed-comments.csv | llm -m m7bi
'It is September 2022 and you are a high-quality
qualitative discourse analyst in the domain of the human
remains trade. This is important for my career.
$input is OCRd text detected from screenshots of social
media posts. There are 303 rows of data. Categorize the nature
of the discourse from each row.

RETURN the ROW NUMBER and the CATEGORIES only.

THOUGHT When I read a new row, I will use the ~ symbol to
indicate a new row.

THOUGHT My categories should capture the nature of the
post or conversation, eg, "EthicalDiscussion" would be useful
if the participants themselves are discussing the ethics of
what they are doing.

THOUGHT There may be more than one appropriate category
for a row.  I should return a maximum of four categories if
appropriate. I will separate categories with a |. I will use
camelCase.

THOUGHT I should try to reuse categories when appropriate.

THOUGHT I should try to categorize by the kind of material
discussed when appropriate.

THOUGHT Participants often adopt an ironic tone, which
will affect the nature of my categorization.

THOUGHT Words like "human" will sometimes be misspelt or
include spaces between the letters to avoid keyword search
detection.' >> output.csv
```

In plain English, 'cat' means 'concatenate', thus the command takes the transcribed-comments.csv file and passes it as input to the LLM command. The -m flag tells the LLM which model to use (here, Mistral-7b-Instruct). The prompt itself is all of the text between the quotation marks. This kind of natural language 'programming' has some important yet curious features; it is not so much about giving clear instructions as it is about steering into the right part of the LLM. First, we give it a date to guide the model towards a time of year where it will have seen *in its training data* high-quality human-written material (a time of year when people might presumably do better work; time of year does seem to make a difference, hence using September 2022, the start of the North American academic year). We give it encouragement to guide the model towards better examples in its training data. We tell it what the data we've just passed to LLM is - screenshots of social media. We tell it what we want ultimately, and then we begin a chain-of-reasoning to further narrow the possibility space of answers. The >> appends the results of the LLM's output for each row of input data to a new row in the output.csv.

Large Language Models are black boxes whose properties are neither well known nor well understood. The prompt above worked very well with GPT-4 as the queried model using short paragraphs of text from Graham's own writings as a test. The much smaller Mistral 7b-Instruct model did not work as well when fed the actual social media posts' transcribed text. Using Mistral 7b-Instruct, 230 unique tags were identified, but when we look at the count of posts by a particular theme, it was apparent that the LLM identified many themes with only one or two examples - overlapping discourses that we observed (as humans) seemed to escape the model's attention. We sampled posts from the various tags identified by the LLM to confirm the accuracy of the model and it seemed that the applied tags were reasonable *in isolation*. But that is not the same thing as saying the tags provided useful insight in terms of actual *discourses*. 'HumanRemains' is not a useful insight in isolation, but it was the tag most often applied by the LLM.

A better experiment might take its cue from methodologies for using agent-based simulation (ABM) in archaeology. In an ABM, one creates a model that captures the elements germane to the phenomena under consideration. Then one runs the model over and over for every combination of possible settings for those elements such that the complete stochastic *landscape* of possible outcomes is understood, comparing the actual data against the combination of parameters that it best matches. The combination of parameters might therefore have something useful to say about the phenomenon (Graham 2024). Similarly, here, one might set the LLM to tag the social media posts over and over again at different settings governing how it generates the likely text (sweeping from the least-probable next tokens to the most-probable across multiple runs etc.); then one might look at the aggregate results in terms of means, medians, and modes. This obviously would be computationally intensive and slow, but perhaps that would not matter if the goal is to characterise materials accurately while respecting the privacy of users.

We will therefore abandon this experiment for the time being. With a stronger underlying language model that could be run on a local machine, we believe that in the future this kind of analysis might provide greater nuance, depth, and insight. It does have the advantage that the entire reading process could be left to the local machine, which would have the virtue of preserving the anonymity of the participants. At present though this approach is hard to replicate and the precise crafting of the prompt or other stochastic elements can lead to unreliable results. There is research to indicate that some LLMs, where the base model is trained from scratch on carefully curated domain-appropriate text, can outperform the models that ingest all available text regardless of its quality or provenance (see e.g. 273 Ventures and Kelvin Legal Data OS 2024). Perhaps a model trained *only* on professional and/or academic archaeological literature could someday be created and serve as the basis for this kind of work (Knibbs 2024; Fairly Trained 2024; Brandsen 2024).

3.3 Visual similarity in photographs accompanying private posts

We can see from the patterns of tagging and the discourses themselves that each group develops its own habits and norms of expression. However, we can see commonalities between not only the main discourses, but the patterning of *interconnected* discourses. We hypothesise that this is a function of the degree of cross-posting that happens between groups. If a person is a member of two different bone-trading groups, they might post the same photographs to each group, and to enable easier posting, use the same text (see Huffer and Graham 2017 on keyword stuffing); the text, however, might be slightly altered to adhere more closely to the norms of the group. But in doing so, the norms of one group might come to influence the norms of a second group.

We might hypothesise that Group 2 and Group 4, for instance, with two broad groupings of similarly interconnected discourses each, might have members or participants in common. In the same way, Group 1 has the most dissimilar patterns of interconnections in its discourses, and so it might have the fewest cross-posts. Because we jettisoned the identifying usernames or handles on posts before we began the analysis, the most obvious vector for investigating this possibility is lost to us (a sensible approach in the future would be to create random codes to stand in for usernames, though not connected to the original usernames, so that authorship remains discernible yet anonymous). However, this causes us to focus more closely on the content and composition of photographs in concert with the discourses of the conversations. Visual similarity of posts can be calculated by pushing the images through an image embedding model. Such models are constructed using neural networks to express the visual features of a photograph as a vector or list of numbers, where each number corresponds to a different dimension of information. In which case, we can express the photographs that accompany these conversations as an image vector and measure for similarity. Images that have almost perfect similarity will be images that have been cross-posted from one group to another. Clusters of

visually similar images might be interpreted as having or responding to similar visual tropes.

There are a variety of ways of calculating visual similarity. The most straightforward is to take an already trained image model and use it to convert the images in question into vectors. We used the Orange Data Mining suite 3.36.1 (Demsar *et al*. 2013), a Python-powered visual programming environment, to create a workflow that ingests our images (in this environment, one 'plugs together' various modules or functions so that inputs and outputs chain together (Figure 5). The workflow file is an xml text encoded file describing the modules and their linkages that can be shared to enable reproducible research (image-workflow.ows). In this case, we used the example image analysis workflow provided in the documentation). The workflow pushes the images through the Inception3 model from Google (which was trained on thousands of examples of photographs and their descriptions). We then visualise the multi-dimensional scaling space of our images to see how the images cluster together.
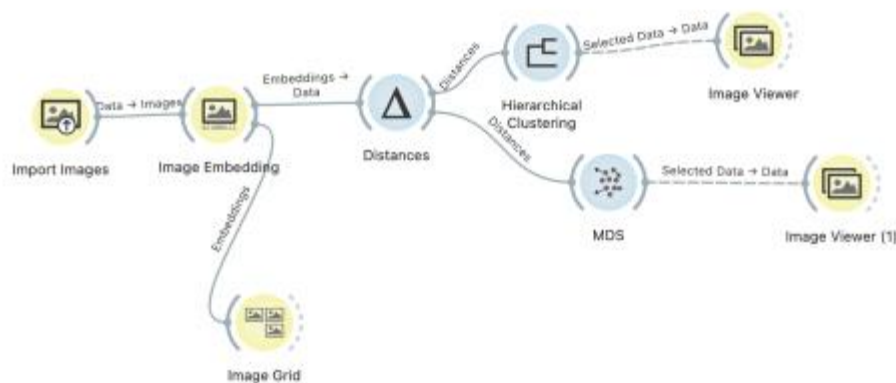

Figure 5: Programming a visual similarity workflow, as represented by the Visual Programming widget on Orange Data Mining
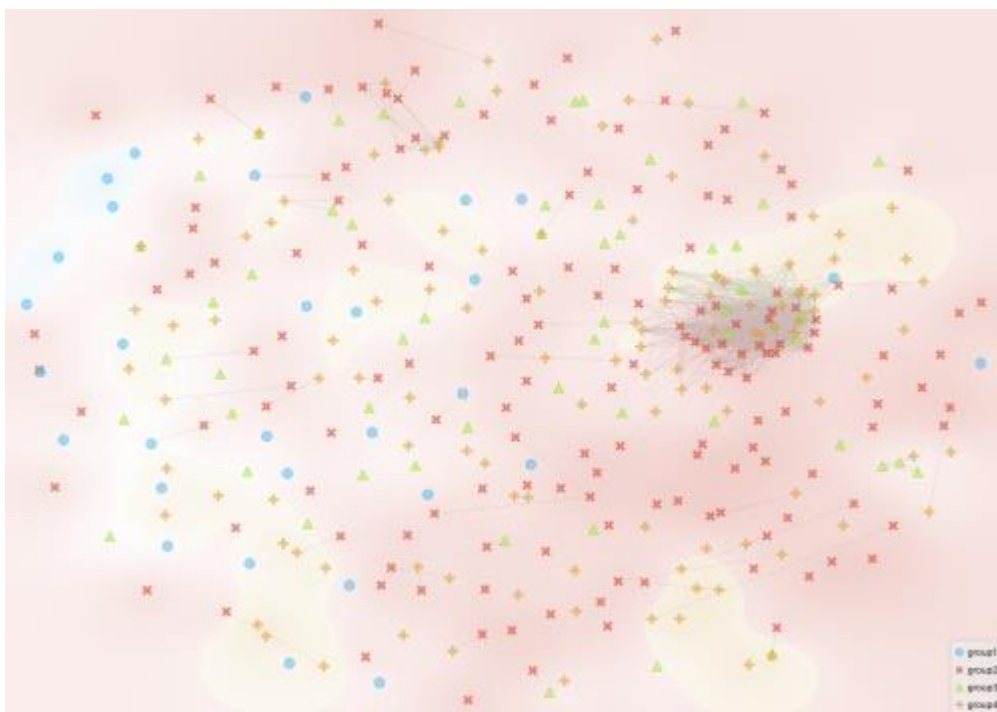
Figure 6: Visual similarity of photographs in posts from the Facebook groups examined in this study, represented via multi-dimensional scaling. Connecting lines represent similarity above a threshold which in this visualisation may be controlled by a slider; the user can then click on data points to manually inspect multiple images

In Figure 6, each point represents an image from the initial post in a conversation across our four groups and within the time frame sampled. We can see quite clearly that while there are images from Group 1 that are cross-posted to other groups, Group 1 is largely isolated from our sample. Group 2 and Group 4, however, are strongly interconnected visually. Indeed, the visual tropes of Group 2 largely set the 'background' in which the human remains trade is expressed in our sample, with small pockets for Group 4 and Group 1 that differentiate. When we re-examine the images from this perspective, we can see that one factor that accounts for this is the professionalised posts of the larger traders: they each have a unique look and feel (in some cases, a mock-gothic aesthetic; in others, a polished storefront; others use the same display table and background across multiple posts).

We might conclude, therefore, that a larger sampling of private Facebook groups connected with the trade in human remains could be differentiated on the visual patterning of their posts alone. We might further imagine that individuals who belong to multiple groups might do so because of their physical location or relationship to potential markets. Thus a mapping of multiple groups' visual similarities could reveal important data about sources of human remains (see Al-Azm and Paul 2019). Alternatively, this patterning might imply a desire to join groups of prominent dealers who might manage several, or a desire to maximise global market access. If one was interested in disrupting the trade, such data would also imply which groups might function as important nodes whose removal (the deletion of the group by Facebook) would have the greatest impact, if even only temporary until such time as the admins of the deleted group formed a replacement. Another possibility is that this visual

similarity data is hinting at the action of Facebook itself as a kind of 'dealer', functioning as designed to bring interested people together: the groups we chose for study are among those that Facebook returns as most germane to our original search query, and that are easiest to join, which illuminates how interest and participation in this trade continues to grow year over year (Graham and Huffer [2020]).

4. Conclusion

This article attempted to map the patterns of discourse within and across conversations and groups pertinent to the human remains trade on Facebook at the end of 2023 via the examination of a large (n=303) number of posts made over a short period of time. We found that private Facebook groups act as a venue for more overt discourses and networks that generate activity within the illicit trade in human remains. Buying and selling, as well as visual aesthetics and subject-centred discussions (including those signalling 'quality sellers' or the dehumanisation of human remains through joking), are generated on these platforms. Across the four groups studied, there is a through-line that we identified which suggests a fair amount of continuity, connecting participants (particularly vendors) in this trade through different venues even as they speak to different virtual audiences comprised of individuals who, off-line, might be local or on the other side of the world.

Within these groups, we observe a clear hierarchy based on the level of professionalisation. 'Professionals', 'occasional sellers', and 'enthusiasts' each have distinct approaches to framing their materials for sale. Groups can be characterised by the relative degree of professionalisation, and groups are connected by individuals who are active across venues, providing vectors for knowledge of how to act in this global, mostly virtual, community so as to spread and normalise the illicit trade in human remains. Photographic evidence alone, when explored through visual similarity, can be used to map these connections, but examining the text of the posts reveals the more overt discourses operating within the closed doors of semi-private Facebook groups.

While prices are openly discussed there is a degree of obfuscation depending on the level of professionalisation, signalling a possible degree of unease with these activities, usually by more 'hobbyist' collectors. The more detailed the backstory, the higher the price commanded. The desire for human skulls and infra-cranial remains, preferably without ornamentation, stands out, while (usually) smaller bones or otherwise discarded fragments find their way into artistic displays, artwork, or jewellery. A striking feature of these conversations is the participants' perceived knowledge of the law (whether correct or not), and how the conversations are laced with formulaic incantations to ward off legal concerns. The disclosure of sellers' general (but sometimes quite specific) locations, possibly to provide shipping estimates, exposes a level of trust within the group but also opens participants up to potential legal risks. Part of the generation of trust lies in the use of in-jokes and 'dad' jokes, which serve to create a sense of camaraderie and an 'edgy' or 'ironic' tone. The public facade of bone traders, characterised by a professed love for the past and a

desire to educate, contrasts with the private gratification found in owning 'something pretty'.

Finally, the ability of bone traders across all levels of professionalisation to navigate the technical intricacies of various digital payment platforms is something not observed often in bone trade posts made in public; such interactions or discussions always take place in private venues such as these groups. They demonstrate a sophisticated understanding of the tools offered by services such as PayPal, Shopify, Venmo, and Meta, leveraging them to facilitate their transactions. Sellers, and especially group administrators, also routinely warn others when things go wrong with payments, or the potential for scams to occur using certain payment systems. In fact, the question of scams and the self-policing of alleged scammers within the human remains trade is a topic well worth further investigation using the methods described above.

A future study of a larger number of groups, which may someday be facilitated by approaches similar to what was attempted here using the Large Language Model, may permit more real-time mapping of the human remains trade and thus increased clarity regarding how human remains cross jurisdictional boundaries. A study of a larger number of groups would also permit mapping the interconnections between groups, pinning key actors to locations and jurisdictions when this isn't already clear from public posts on Facebook and elsewhere (much like the ATHAR Project's 2019 study of antiquities trading on Facebook was able to achieve). As we found through this study, a close reading of these discourses is, at present, best done by humans, and may be facilitated on a large scale by OCR technology.

Acknowledgements

Data Availability

The source data from social media posts is not available, as per our research protocol. The image analysis workflow for the Orange Data Mining platform is available as the file 'image-workflow.ows'. The configuration of the Obsidian text-editing platform for the performance of the qualitative analysis is the same as that provided by Murphy 2021. Object Character Recognition was performed using the open source 'PaddleOCR' from PaddlePaddle 2023.

Footnotes

1. See Graham 2023. Obsidian is available for multiple platforms. As of the time of writing, Meta appears to be creating an API for approved researchers to obtain aggregated data from the Facebook platform. Facebook's current terms of service prohibit automated scraping, though with caveats, see https://www.facebook.com/apps/site_scraping_tos_terms.php. There appears to be no prohibition in Facebook's terms of service against collecting screenshots. ←

2. Copying and pasting posts directly into software like Word or Excel would preserve the text of posts, but in an unstructured way, and would complicate working with the associated images for each post. The present workflow would also permit easier collocation of image data (like object detection) with the associated texts and conversations. The OCR package used was PaddleOCR 2.7.1. We tested other open-source OCR packages and found that PaddleOCR gave best results detecting the mixture of clean text and text within the original post's associated photographs, such as with overlaid text (e.g. as done with MS Paint or PowerPoint), emojis, and other textual elements. ←

3. This was not an experiment that could be run ethically with something like GPT4, the current 'best' LLM available. GPT4 requires vast computing and energy resources to process data, so to use it one would have to send the data to OpenAI's computers; it is not certain what happens to the data once sent to OpenAI's servers or whether it would get used for further training of OpenAI's models. ←

Bibliography

Al-Azm, A. and Paul, K.A. 2019 *Facebook's Black Market in Antiquities: Trafficking, Terrorism, and War Crimes*, ATHAR Project. https://static1.squarespace.com/static/5e3a7fb845f8c668df48d437/t/5e4e58ef0caa64606351a54b/1582192944742/ATHAR-FB-Report-June-2019.pdf

Baumer, E.P.S., Mimno, D., Guha, S., Quan, E. and Gay, G.K. 2017 'Comparing grounded theory and topic modeling: extreme divergence or unlikely convergence?, *Journal of the Association for Information Science and Technology* **68**(6), 1397-1410. https://doi.org/10.1002/asi.23786

Brandsen, A. 2024 'Archaeology specific BERT models for English, German, and Dutch (Version 4)', Paper given at Computer Applications and Quantitative Methods in Archaeology Conference 2023 (CAA2023), Amsterdam. https://doi.org/10.5281/zenodo.10610882

Breda, E. 2023 *Bone Collectors: Personhood and Appeal in the Sale and Trade of Human Skeletal Remains on Facebook*, Unpublished Masters thesis, Florida Atlantic University, Boca Raton, Florida.

Charmaz, K. 2017 'The power of constructivist grounded theory for critical inquiry', *Qualitative Inquiry* **23**(1), 34-45. https://doi.org/10.1177/1077800416657105

Davidson, K., Graham, S. and Huffer, D. 2021 'Exploring taste formation and performance in the illicit trade of human remains on Instagram' in N. Oosterman and D. Yates (eds) *Crime and Art: Sociological and Criminological Perspectives of Crimes in the*

Art World, Studies in Art, Heritage, Law and the Market, Cham: Springer International Publishing. 29-44.

Davidson, K., Graham, S. and Huffer, D. 2024 'A protocol for when social media goes private: studying archaeological or heritage discourses in closed Facebook groups', *Internet Archaeology* **67**. https://doi.org/10.11141/ia.67.11

Demsar, J., Curk, T., Erjavec, A., Gorup, C., Hocevar, T., Milutinovic, M., Mozina, M., Polajnar, M., Toplak, M., Staric, A., Stajdohar, M., Umek, L., Zagar, L., Zbontar, J., Zitnik, M. and Zupan, B. 2013 'Orange: Data Mining Toolbox in Python', *Journal of Machine Learning Research* **14**, 2349-2353. http://jmlr.org/papers/v14/demsar13a.html

Dillon, D.R. 2012 'Grounded theory and qualitative research' in C.A. Chapelle (ed) *The Encyclopedia of Applied Linguistics*, 1st edition, Wiley. https://doi.org/10.1002/9781405198431.wbeal0486

Dundler, L. 2021 '#antiquitiesdealers' in N. Oosterman and D. Yates (eds) *Crime and Art: Sociological and Criminological Perspectives of Crimes in the Art World*, Studies in Art, Heritage, Law and the Market, Cham: Springer International Publishing. 45-61.

Fairly Trained 2024 'Announcing our newly certified models (including LLM & speech) and new supporters', *Fairly Trained* [blog], 20 March 2024, https://www.fairlytrained.org/blog/newly-certified-models-march-24

Google 2024 'Advanced Guide to Inception 3', *Google Cloud*. https://cloud.google.com/tpu/docs/inception-v3-advanced

Graham, S. 2023 'Qualitative analysis of social media posts - a workflow from OCR to Obsidian', *X-Lab* [blog], 15 December 2023. https://carleton.ca/xlab/2023/qualitative-analysis-of-social-media-posts-a-workflow-from-ocr-to-obsidian/

Graham, S. 2024 'The behaviour space(s) of Large Language Models', *X-Lab* [blog], 9 April 2024. https://carleton.ca/xlab/2024/the-behaviour-spaces-of-large-language-models/

Graham, S. and Huffer, D. 2020 'Reproducibility, replicability, and revisiting the Insta-Dead and the human remains trade', *Internet Archaeology* **55**. https://doi.org/10.11141/ia.55.11

Graham, S., Huffer, D. and Blackadar, J. 2020a 'Towards a digital sensorial archaeology as an experiment in distant viewing of the trade in human remains on Instagram', *Heritage 2020* **3**(2), Art and Antiquities Crime special issue, 208-27. https://doi.org/10.3390/heritage3020013

Graham, S., Lane, A., Huffer, D. and Angourakis, A. 2020b 'Towards a method for discerning sources of supply within the human remains trade via patterns of visual

dissimilarity and computer vision', *Journal of Computer Applications in Archaeology* **3**(1), 253-68. https://doi.org/10.5334/jcaa.59

Graham, S., Huffer, D. and Simons, J. 2022 'When TikTok discovered the human remains trade: a case study', *Open Archaeology* **8**(1), 196-219. https://doi.org/10.1515/opar-2022-0235

Graham, S., Yates, D. and El-Roby, A. 2023 'Investigating antiquities trafficking with Generative Pre-Trained Transformer (GPT)-3 enabled knowledge graphs: a case study', *Open Research Europe* **3**, 100. https://doi.org/10.12688/openreseurope.16003.1

Huffer, D. and Graham, S. 2017 'The Insta-Dead: the rhetoric of the human remains trade on Instagram', *Internet Archaeology* **45**. https://doi.org/10.11141/ia.45.5

Huffer, D. and Graham, S. 2018 'Fleshing out the bones: studying the human remains trade with Tensorflow and Inception', *Journal of Computer Applications in Archaeology* **1**(1). https://doi.org/10.5334/jcaa.8.

Huffer, D. and Graham, S. 2023 *These Were People Once: The Online Trade In Human Remains, and Why It Matters*, New York: Berghahn.

Huffer, D., Wood, C. and Graham, S. 2019 'What the machine saw: some questions on the ethics of computer vision and machine learning to investigate human remains trafficking', *Internet Archaeology* **52**. https://doi.org/10.11141/ia.52.5

Huffer, D., Guerreiro, A., and Graham, S. 2021 'Osteological assessment of a seized shipment of modified human crania: Implications for Dayak cultural heritage preservation and the global human remains trade', *Journal of Borneo-Kalimantan* **7**(1), 67-93.

Huffer, D., Simons, J., Brughmans, T. and Graham, S. 2022 'Alleen voor studiedoeleinden (For study purposes only): The human remains trade on Marktplaats.nl', *Anthropologica et Praehistorica* **131**, 179-94. https://doi.org/10.57937/ap.2022.008

Knibbs, K. 2024 'Here's proof you can train an AI Model without slurping copyrighted content', *Wired*, 20 March 2024. https://www.wired.com/story/proof-you-can-train-ai-without-slurping-copyrighted-content/

273 Ventures and Kelvin Legal Data OS 2024 'Kl3m.Ai - the cleanest LLM in the world', *KL3M*, https://kl3m.ai

Mistral AI, 2023 'Announcing Mistral 7B The best 7B model to date, Apache 2.0', *Mistral AI*, 27 September 2023. https://mistral.ai/news/announcing-mistral-7b/

Moore, J., Magee, S., Gamreklidze, E. and Kowalewski, J. 2019 'Social media mourning: using grounded theory to explore how people grieve on social networking

sites', *OMEGA - Journal of Death and Dying* **79**(3), 231-59. https://doi.org/10.1177/0030222817709691

Murphy, R.J.A. 2021 'An integrated qualitative analysis environment with Obsidian', *Fulcra Design* [Blog], 13 August 2021. https://fulcra.design/Posts/An-Integrated-Qualitative-Analysis-Environment-with-Obsidian/

Nelson, L.K. 2020 'Computational grounded theory: a methodological framework', *Sociological Methods and Research* **49**(1), 3-42. https://doi.org/10.1177/0049124117729703

Tri-Council: Canadian Institutes of Health Research, Natural Sciences and Engineering Research Council of Canada, and Social Sciences and Humanities Research Council of Canada 2022 *Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans.* https://ethics.gc.ca/eng/policy-politique_tcps2-eptc2_2022.html

Willison, S. 2023 'LLM A CLI utility and Python library for interacting with Large Language Models, both via remote APIs and models that can be installed and run on your own machine'. https://llm.datasette.io/en/stable/ Version 0.13.1