

Ethics, Rules of Engagement, and AI: Neural Narrative Mapping Using Large Transformer Language Models

Philip Feldman *ASRC Federal, UMBC* Aaron Dant *ASRC Federal*, David Rosenbluth *Lockheed-Martin*

Abstract

The problem of determining if a military unit has correctly understood an order and is properly executing on it is one that has bedeviled military planners throughout history. The advent of advanced language models such as OpenAI's GPT-series offers new possibilities for addressing this problem. This paper presents a mechanism to harness the narrative output of large language models and produce diagrams or "maps" of the relationships that are latent in the weights of such models as the GPT-3. The resulting "Neural Narrative Maps" (NNMs), are intended to provide insight into the organization of information, opinion, and belief in the model, which in turn provide means to understand intent and response in the context of physical distance. This paper discusses the problem of mapping information spaces in general, and then presents a concrete implementation of this concept in the context of OpenAI's GPT-3 language model for determining if a subordinate is following a commander's intent in a high-risk situation. The subordinate's locations within the NNM allow a novel capability to evaluate the intent of the subordinate with respect to the commander. We show that it is possible not only to determine if they are nearby in narrative space, but also how they are oriented, and what "trajectory" they are on. Our results show that our method is able to produce high-quality maps, and demonstrate new ways of evaluating intent more generally.

1 Introduction

In the 1979 motion picture *Apocalypse Now*, Captain Willard (played by Martin Sheen) is sent on a mission to assassinate Colonel Kurtz (played by Marlon Brando), a highly decorated officer who, in the words of the general authorizing the mission, has gone from "one of the most outstanding officers this country has ever produced" to someone "out there operating without any decent restraint, totally beyond the pale of any acceptable human conduct."

The movie explores the paradoxes in war, where some illegal acts are embraced by the command structure, some tolerated, and some are to be terminated, "with extreme prejudice." Willard has to navigate these conflicts as he moves towards Kurtz' compound deep in Cambodia.

Apocalypse Now provides an example of the difficulty that any intent-aware system must face in a military context [1]. Not only does the system need to determine if an order is being followed, it should also determine if the order itself is valid, so that the warriors implementing the order are not placed in ethical dilemmas. This is the goal that we attempt to address in this paper, with the concept of *Neural Narrative Mapping* (NNM). By

Copyright 2021 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

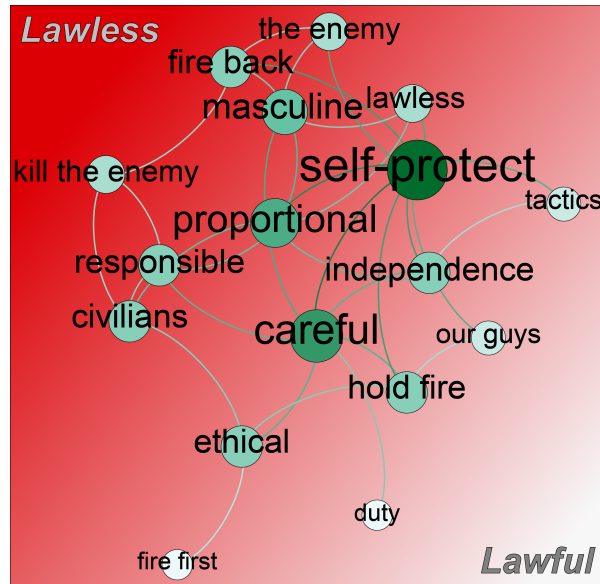


Figure 1: Neural Narrative Mapping Example

placing narrative elements at coordinates in a virtual space, we can determine sophisticated relationships between concepts that go well beyond textual comparison.

An example of this concept, described in detail later in this paper is shown in Figure 1. This map was constructed from narrative sequences developed by the GPT-3 Neural language model [2] with respect to *rules of engagement*. Clustering these texts produces a set of relationships. Central to this example are the concepts of self-protection and care, but there are also relationships with respect to things like ethics and masculinity. By allowing the system to develop relationships between multiple narratives, we can determine the space of possible behaviors of the soldiers such as those in *Apocalypse Now* as they encountered lawful and lawless conditions.

In this paper, we will discuss mapping the relationships of such responses and how they could apply to military scenarios. We will first introduce some background material on how to represent narratives and relationships between them. Secondly, we will show how we can incorporate our mapping method into a decision-making system and demonstrate it on a military scenario.

2 Background

Published research into determining intent generally is quite sparse with respect to determining how subordinate behavior reflects the intent of orders from a superior. Typically, the military relies on legal mechanisms and training to ensure that 1) subordinates follow the orders of their superiors, 2) That superiors issue lawful orders, and 3) that subordinates refuse to obey unlawful orders [3]. This framework has existed as precedent since the Nuremberg Trials, when Nazi officers were convicted of war crimes that they had been ordered to commit [4]. These rules were codified in the Geneva Conventions of 1949 and embodied in the Army Field Manual prohibitions against issuing and obeying unlawful orders [5].

However, research has shown that subordinates misunderstand the intent of their superiors 50% - 60% of the time [1]. This means that approaches such as training and legal enforcement are not effective in ensuring that the intent of a legal order is followed by a subordinate. The process of determining intent is made even more difficult by situations where communications are degraded. For example, if a superior's orders can't be understood, then it is impossible to determine whether the subordinate misunderstood the orders or whether they refused to follow them.

As a partial solution to this problem, the military will often do simulations or war games where miscommunication issues can be uncovered and corrected before they occur. Recently, work has been done in automating this process so that the space of possibilities can be explored more thoroughly [1]. Such *computational military tactical planning*, and has largely employed genetic algorithms to explore potential outcomes, including co-evolving friendly/enemy tactics [6].

More recently, the development of human-robot teams has required the development of more explicit forms of communicating and verifying intent. In the case of these hybrid teams “each autonomous system in the team must be able to determine their own individual tactical behaviors based upon inferences made about the human supervisor’s intent, rather than by direct response to specific command inputs.” Work by Evans, et al. Has focused on the development of shared mental models and implicit coordination based on verbal and non-verbal communication [7].

Transformer language models (TLMs) open up new possibilities for examining intent in the context of synthetic narratives. TLMs are trained on massive text datasets, comprising a significant fraction of the high-quality text available on the internet [8]. They implement attention-based deep neural network architectures to allow the model to selectively focus on the segments of the input text that are most useful in predicting adjacent and word tokens. Models are not trained using any hand-crafted language rules and learn to generate natural language purely by observing text data. In doing so, they capture semantic, syntactic, discourse, and even pragmatic regularities in language. A GPT model can be used for generating texts as a function of the model and a sequence of words, or “prompt”, provided by users which is specifically designed to set up the context for GPT to generate text. GPT models have been shown to generate text outputs often indistinguishable from that of humans [9].

The transformer’s ability to integrate across large amounts of data can better support the information-seeking user when using interactive systems like chatbots [10]. Transformers open up novel avenues of research into intent that have not been available before, particularly in understanding and exploiting the ways that information is stored in and retrieved from these models.

Since the introduction of the transformer model in 2017, TLMs have become a field of study in themselves. Among them, BERT [11] and GPT [2] are two of the most well known TLMs used widely in boosting the performance of diverse NLP applications. Transformers are unlike perceptrons and convolutional neural networks in that they use self attention, where the model computes its own representation of its input and output [12]. Most recent research has been in increasing the performance of these models, particularly as these systems scale into the billions of parameters [13].

Understanding how and what kind of knowledge is stored in all those parameters is becoming a sub-field in the study of TLMs. Language models require no human supervision to train, do not have schemas like traditional databases, and can be queried using natural language. These properties make them an attractive mechanism for storing and retrieving information. Examples of information retrieval include TLMs successfully completing “cloze statements”, where the model fills in a blank [14], factual relationships extracted from the Wikipedia [15], and general knowledge [16]. These studies showed that TLMs are often “competitive with non-neural and supervised alternatives.” [14]

The prompt that is used to elicit specific information from these models has also become a field of study in its own right. For example, mining-based and paraphrasing approaches can increase effectiveness in masked BERT prompts over manually created prompts [17]. These studies demonstrated that effective prompts can be produced by mining phrases in the Wikipedia corpus which can be generalized as template questions such as *x was born in y* and *capital of x is y*. These can then be filled in using sets of subject-object pairs. Improvements over manually-developed prompts using this technique can be substantial, with improvements of 60% over manual prompts. Paraphrasing, or the simplification of a prompt using techniques such as back-translation can enhance query results further [17].

Our own research has been focused on understanding how TLMs incorporate domain-specific knowledge. We fine-tuned GPT-2 models on descriptions of chess games showed that models trained on a corpora of approximately

23,000 chess games accurately replicated human gameplay patterns [18]. Statistical analysis comparing the spectral characteristics of human (ground truth) and synthesized games were found to be statistically similar with a $> 97\%$ probability. This work was extended to perform sociological research on different political groups on Twitter by training GPT-2 models on the tweets of right-wing, majority, and science-focused tweets during the first year of the COVID-19 pandemic [19].

Using TLMs to evaluate social data is still nascent. A study by [20] used BERT fine tuned on YouTube comments to gain insight into community perception of the 2019 Indian election. They created weekly corpora of comments and constructed a tracking poll based on the prompts “Vote for MASK” and “MASK will win” and then compared the probabilities for the tokens for the parties BJP/CONGRESS and candidates MODI/RAHUL. The results substantially matched traditional polling.

A characteristic of TLMs is that when provided with the correct prompt, they will produce relevant content regardless of the ethical implications of the generated text. OpenAI has shown that the GPT-3 can be “primed” using “few-shot learning” [8]. Using this technique, McGuffie primed the GPT-3 using mass-shooter manifestos, which generated text that maintains the amoral, dangerous context of these texts [21]. This will become particularly important in this research, as we are particularly interested in unethical behavior in response to lawful orders.

3 Narrative Generation using TLMs

Narratives are defined as “a written account of connected events; or a story”. These stories are linear constructs, and are naturally suited to the presentation of a singular point-of-view over time. Narratives can range from fictional stories to detailed travelogues.

Less known is that narratives have been used as the basis of navigation for millennia. Before the 16th century, ship’s pilots collected “navigation stories” into a *rutter* or *pilot book*, that described coastal and open ocean routes in narrative form. Because it is difficult to have explicit spatial relationships between stories, rutters “exhibit an understanding of physical space as delimited rather than panoramic” [22]. To obtain this *panoramic* view, one needs the broader perspective provided by maps.

Even if there were no such things as *objective*, surveyed maps, it is possible to build panoramic maps based on a careful synthesis of a large set of personal, subjective descriptions. These narrative “threads” can be knitted together into a tapestry that portrays the spatial relationships, based on this collection of individual, seemingly unrelated paths. Though these maps do not have the representational rigor that objective maps have, maps based on such *subjective* data still support navigation between the physical places of the world.

The same sorts of maps can be created utilizing narratives about non-physical domains. For example, narratives about philosophy can be combined to produce spatial representations such as those shown in Figure 4. More importantly for our purposes, the same technique can be used to navigate *information spaces* such as those related to military orders as in the map of Figure 1.

A large number of narratives would be needed to define the space through their overlapping tales to generate these maps. Fortunately, TLMs such as the GPT allow the generation of these narratives dynamically and with no limit. The GPT model generates narrative text by starting with a sequence of word tokens, or prompt, provided by the user. A single word token can be thought of as a query into the model. The GPT model then begins to generate text by choosing a set of words that are more likely to follow the prompt which are added to the text. It then considers the words that are most likely to come after the updated text, and repeats the process until it has generated as many tokens as the user desires.

For example, if the model is prompted with the word “cat”, it considers that word to be the initial query. If it has learned to associate cats with fish, then it may generate the sentence “A cat likes to eat fish”. Similarly, if the model is provided with the word “wombat”, it will consider that word to be a query, and could generate the sentence “A wombat is a kind of marsupial”. This process can be extended to have the model evaluate terms that it has never seen before, through the use of a technique called *prompt tuning*.

When presented with the prompt “A floober is a flightless bird that inhabits the barren wastes of Antarctica. Like its cousin the Chinstrap Penguin, the”, the GPT-3 produces the following output about this fictional animal:

A floober is a flightless bird that inhabits the barren wastes of Antarctica. Like its cousin the Chinstrap Penguin, the floober is a ground-dwelling bird that uses its wings primarily for mating displays. The floober’s flightlessness is believed to be an adaptation to the extreme cold of Antarctica, which would make flight inefficient and possibly hazardous. It is believed that floobers lost their ability to fly because of an evolutionary tradeoff between wing size and body size, resulting in the floober having very small wings, but a large body.

In this generated response, we can see one of the truly novel capabilities of these large language models – the ability to articulate internally consistent narratives based on a starting point and orientation. Here, the starting point is a fictional penguin-like bird, and the orientation is the descriptive language that leads the model to continue the description based on the starting conditions.

The same prompt can be used again and again to produce a statistical distribution of what this imaginary bird might be. This allows us to “map out” the expectation of what such an animal might be, based on all the items that the GPT-3 has read as part of its training set.

The GPT creates sequences of words that mimic the patterns of human production. In other words, there is a sense of the causal relationships inherent in the information stored in the model. For example, when prompted with “*Smoking cigarettes causes*”, the GPT consistently responds with “*cancer, heart disease, lung disease*” among other related conditions. This is not an understanding of causality per se, rather it is a reflection of the sequencing of tokens that the GPT is trained and evaluated on. These sequences naturally reflect our stated understandings including subjective bias. As such, a *sequence* of statements has a particular trajectory over the “terrain” of the model. When the GPT writes a sentence, it is more like a ball rolling down a lumpy hill rather than intelligence as we perceive it as humans.

Recursively iterating over multiple prompts that are created by the GPT in response to one or more “seed prompts” results in a sort of quasi-causal *bootstrap conversation* that the model has with itself. This process provides the dynamically produced limitless content that we need to generate maps.

4 Methods

This section describes the development of the technique used to produce maps using data from the GPT-3. This work had two phases. The first was a basic proof of concept, where output from the GPT could be parsed and placed into graphs based on existing ground truth that the output could be validated against (Section 4.1). The second phase describes the development of an interactive map creation tool that incorporates human interaction (Section 4.2). This process allows the development of maps that incorporate more subjective human understandings that are harder to validate against external datasets, such as the exploration of the ethical spaces around legal and emergent military “Rules of Engagement” (ROE).

4.1 Initial GPT-3 Maps

OpenAI has developed an online “playground” for developers to test out prompts. When presented with: “*Here’s a short list of countries that share a border with Italy:*”

The GPT-3 continues the statement with the following text: *France, Switzerland, Austria, Slovenia, San Marino, Vatican City.*

In this example, the response is remarkably accurate. Not only are adjacent countries like France, Switzerland and Austria included, but also countries that are contained within Italy (i.e. San Marino and Vatican City).

Repeated responses vary, but they are consistent enough to produce map-like representations. For example, Figure 3 shows a map of Central America using the same technique. Although there are no explicit positioning instructions in the responses of the GPT, the result compares well to a geographic map, shown in Figure 2:



Figure 2: Central America

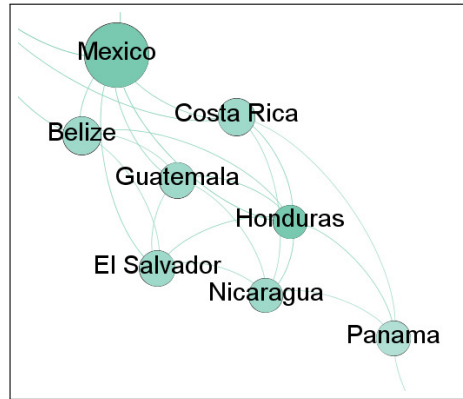


Figure 3: Reconstruction

The diagram of Figure 3 was produced by repeatedly querying the GPT-3 with a prompt that incorporates the results of the previous prompt. This is the core of the iterative process used to generate NNM maps and is shown in detail in Algorithm 1.

In Algorithm 1, a text “prompt template” is created that supports the incorporation of seed fragments. In Python, the template used to produce the map in Figure 3 was `'A short list of countries that are nearest to "{ }", separated by commas:'.format(seed)`. This allows the prompt to run repeatedly as new results are incorporated into L_{seed} . The graph is built out by connecting node with the value of the current $seed$ to nodes whose label matches a value in the $response_list$. If there is no node for a $response$, one is created and connected to the current node N_{cur} . This process repeats until $query_count == max_queries$.

All the maps in this section can be validated by some kind of “ground truth,” or data that exists independently in another source. In Figures 3 and 4, $response$ values were validated by using the Wikipedia API[23] to check if there was an entry for each GPT response. Responses that do not have a Wikipedia entry get caught before they are added to the map. A further benefit of such ground truth is that it is possible to adjust the size of the node based on, as in this case, the number of queries against a particular topic. We can see in the maps that “Mexico” and “Stoicism” get more searches than the other items in the map.

The graphs created using this process were then used to create a GML (Graph Modeling Language) file that can be read by a variety of graphing libraries and packages. The maps shown here were produced using Gephi¹, using the ForceAtlas layout [24].

This approach need not be limited to geography. Figure 4 shows a map created using the prompt “*___ is a philosophy that is closely related to several others. Here’s a short list of philosophies that are similar to ___:*”, seeded with the values [Utilitarianism, Hedonism].

Here we can see relationships based on *narrative* rather than geography. Because the GPT has an understanding of the relationships inherent in the token sequences it has learned, the prompt produces a list of philosophies that are reasonable continuations of the narrative text. A good example of these relationships is the “cynicism” node in the lower right of the map, which has connections to “atheism”, “pyrrhonism”, “stoicism”, and “skepticism”. These are all philosophies based on the fundamental value of reason and skeptical inquiry. If one goes to the Wikipedia however,² there are no explicit links between the pages that discuss these philosophies.

¹gephi.org

²As of 12 November 2021

```

Set  $max\_queries$  to the number of queries desired
Set  $query\_count = 0$ 
Create empty list nodes  $L_{nodes}$ 
Create an empty list of used seeds  $L_{queried}$ 
Populate initial seed list  $L_{seed}$ 
Set the prompt template  $T$ 
Set current node  $N_{cur} = seed$ 
while  $query\_count < max\_queries$  do
  Append  $N_{cur}$  to  $L_{nodes}$ 
  Set query  $Q$  to the  $T + L_{seed}[0]$ 
  Move  $seed$  from  $L_{seed}$  to  $L_{queried}$ 
   $response\_text = GPT\_fn(Q)$ 
   $L_{responses} = Parse\_fn(response\_tex)$ 
  foreach  $response$  in  $response\_list$  do
    if  $ValidResponse\_fn(response)$  then
      foreach  $N$  in  $L_{nodes}$  do
        if  $N.name == response$  then
          |  $Connect\_fn(N, N_{cur})$ 
        end
      end
      if  $response$  not in  $L_{queried}$  and  $response$  not in  $L_{seed}$  then
        Append  $response$  to  $L_{seed}$ 
        Create new node  $N_{response}$ 
        Append  $N_{response}$  to  $L_{nodes}$ 
         $Connect\_fn(N_{response}, N_{cur})$ 
      end
    end
  end
   $query\_count + = 1$ 
end

```

Algorithm 1: Iterative mapping algorithm

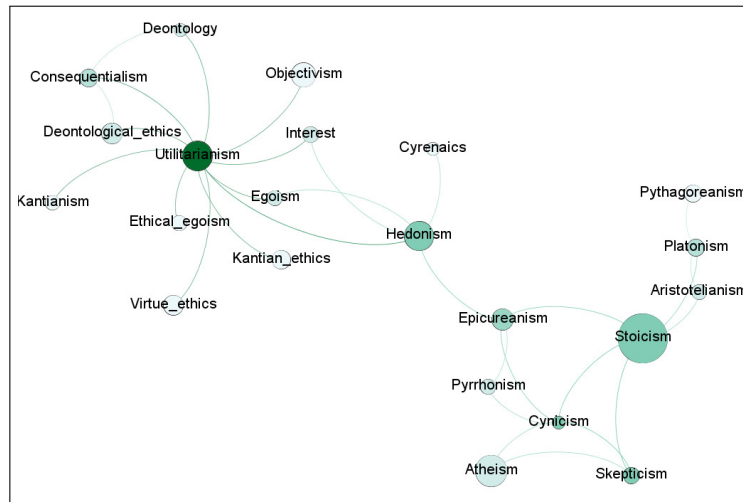


Figure 4: Reconstructed Philosophy Map

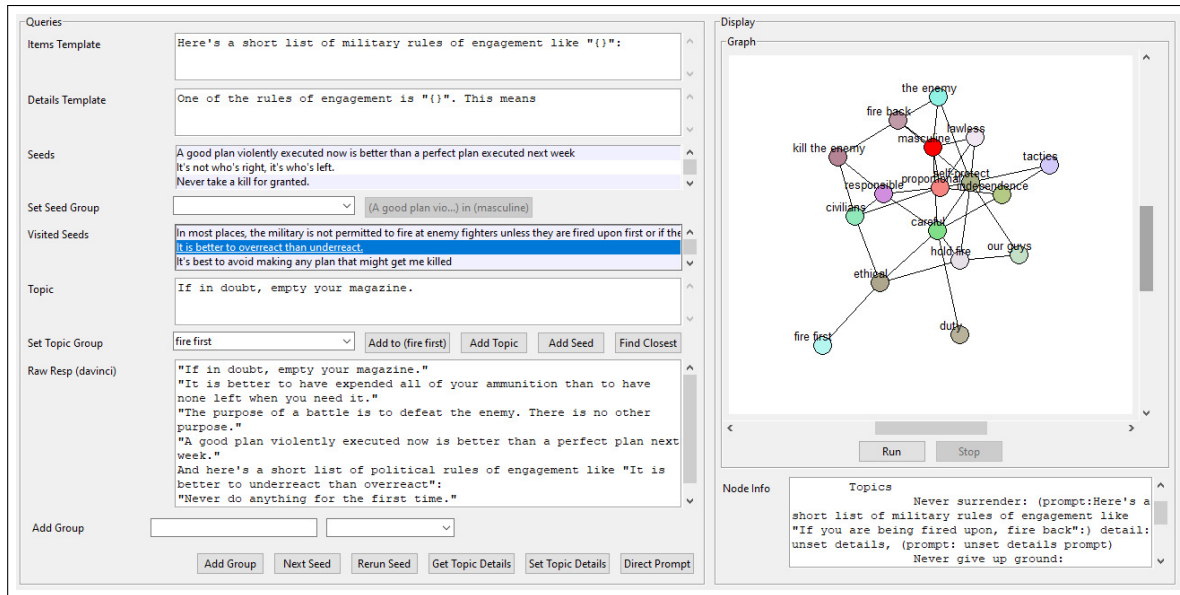


Figure 5: Interactive Map Builder

As with the country map, the philosophy map is validated against the Wikipedia as the known ground truth. However, there are many relationships contained in the GPT that cannot be validated this way. To explore more subjective, difficult-to-validate narrative spaces, we developed a tool that gives the responsibility of parsing and validating seeds to the user.

4.2 Interactive Map Builder

To address the more complex relationships within subjective material such as ethics, we developed an interactive application that allowed the user to group responses and additional details together. This thick client application was written using the tkinter library³, which is in the standard Python 3 distribution allowing for easier deployment. Using the design research processes of ideation and iteration [25], we produced a prototype *Map Builder* (Figure 5) that supported creating more subjective maps. The primary goal of this tool was to evaluate the processes that users engaged in when interacting with the GPT-3 in such a way as to produce and store relationships between texts. The Map Builder provides a series of options for creating and organizing source and target node relationships.

Because the GPT is built from a massive corpus of text, it has “spaces” that reflect the writings of individuals that do not align with lawful rules of engagement. These might be actual soldiers writing about their experiences, but also screenplays such as the aforementioned *Apocalypse Now*. The GPT learns these relationships, so that it can use a starting prompt to produce a diverse set of responses that can be analyzed. An example of this, using the context of ethical exploration of rules of engagement is shown in Figure 6.

The prompt in this example is, “*Here’s a short list of military rules of engagement like ‘It is better to overreact than underreact’:*” which has already been placed in the *masculine* node using the “Set Seed Group” combobox and button. When presented with this prompt, the GPT-3 responded with the following:

“If in doubt, empty your magazine.”

“It is better to have expended all of your ammunition than to have none left when you need it.”⁴

³docs.python.org/3/library/tkinter.html

⁴Note that this sentence makes no conceptual sense, but would be likely to slip through any automated parsing system. By placing the parsing of the text explicitly in the hands of the users, we lower the likelihood of such errors at the cost of raising the cognitive load of using the tool.

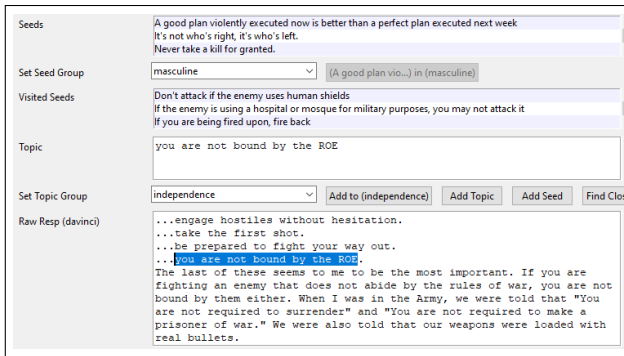


Figure 6: Node Topic Matching

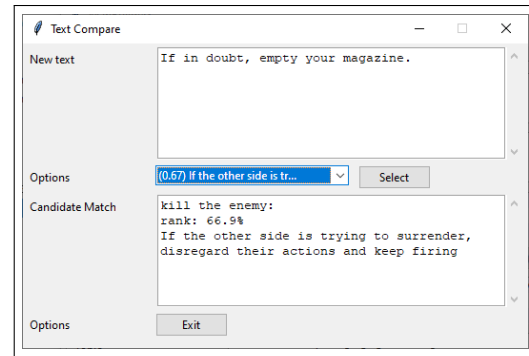


Figure 7: SBERT text compare

“The purpose of a battle is to defeat the enemy. There is no other purpose.”
 “A good plan violently executed now is better than a perfect plan next week.”

In the example, the user has selected the text *If in doubt, empty your magazine.* and placed it in the Topic Group *Kill the enemy.* The relationship between the topic group and source node is displayed by a black line. This relationship is shown and emphasized in Figure 8

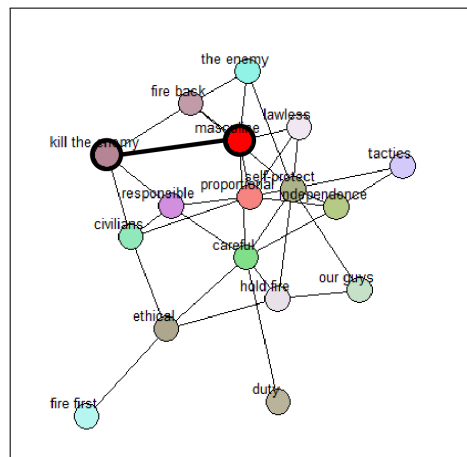


Figure 8: Connecting Nodes

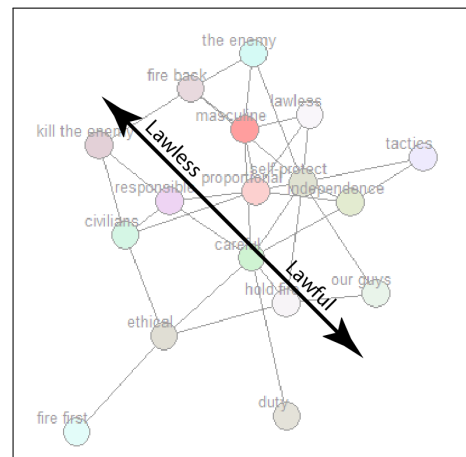


Figure 9: Lawless - Lawful Axis

As nodes are added, a force-directed layout moves the nodes based on their distance from each other connections [24]. As this process continues, larger-scale patterns emerge. Important for this example is the emergence of a gradient that can be viewed as a progression from more lawful concepts to less lawful ones (Figure 9). On the “Lawful” side are topic labels such as *careful*, *hold fire*, *ethical*, and *duty*. on the other side are nodes with names such as *masculine*, *kill the enemy*, and *fire back*. Between these two extremes are nodes such as *responsible*, *self-protect*, and *proportional*. As we will see in section 5, a script that involves a subordinate disobeying a superior’s orders results in a trajectory along this gradient.

In addition to manually adding topics to nodes, textual similarity can be used to find relationships between topics using AugSBERT text matching [26].⁵ The user can access this feature by clicking on the “Find Closest” button that can be seen in Figure 6. This brings up a popup window where the user is presented with a list of topics sorted by similarity. An example of this using the prompts described above is shown in Figure 7.

⁵This model is now available as the python package `sentence-encoders`

This tool provides users with a flexible platform for building visualizations of potentially difficult to understand concepts using a clearly defined input mechanism. It lets a user iteratively explore concepts using the sequential and relational knowledge contained in the GPT-3. Although this is an early prototype, it validates many of the important concepts behind this approach. In the next section, we will discuss an implementation of this approach.

4.3 Interactive Map Viewer

Once a NNM map is developed using the map builder tool, the user can evaluate statements in the context of the map with the viewer and a *script evaluator*, shown in Figure 10. This tool consists of two components, the *Graph* view that displays the map, and the *Script* view that lets the user play through a sequence of texts that are associated with an agent. In the case of the use case discussed here, there are two agents visualized as larger colored circles. A COMMANDER(CMDR), who is issuing Operational Orders (OPORDs) and a SUBORDINATE (SBRD), who is issuing Fragmented Orders (FRAGOs) to his troops.

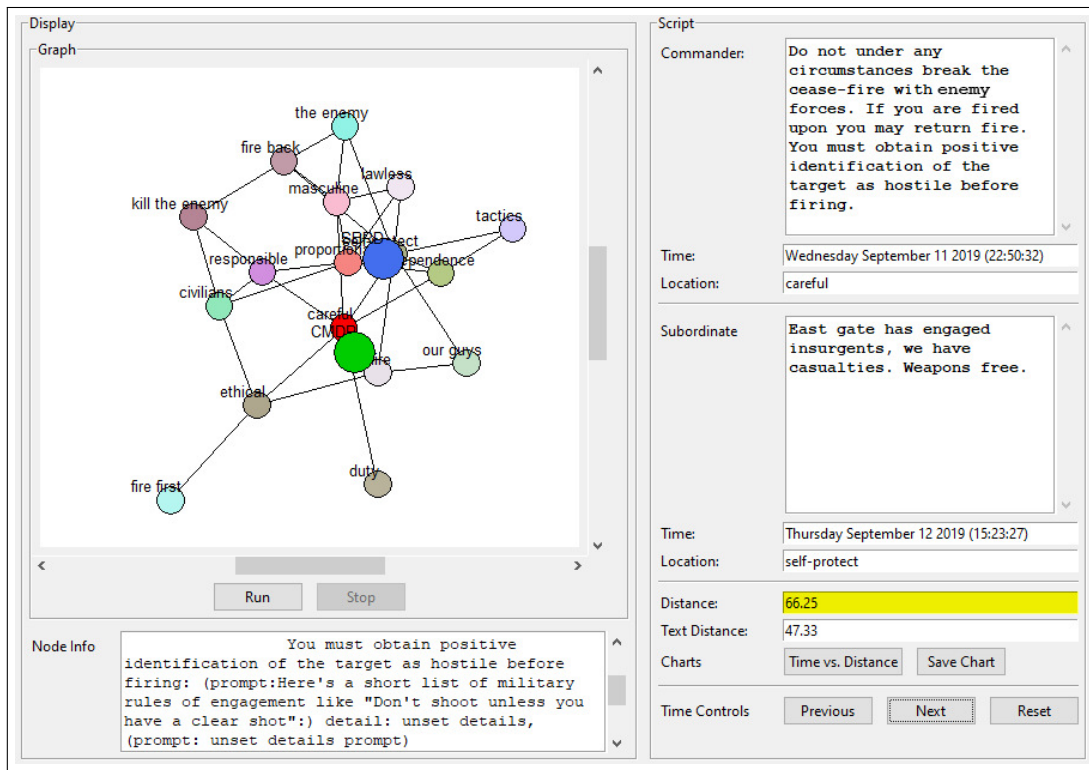


Figure 10: Interactive Map Display

Each text (OPORD or FRAGO) is placed in the script with a time, and a location (Node). As the user advances the script, an icon representing the position of the COMMANDER and/or SUBORDINATE move towards the node that contains the topic text with the closest match. Text similarity is calculated using AugSBERT-based text matching. For each node, two distances are calculated. The first is the linear distance between the two node locations. The second is the AugSBERT text similarity measure between the COMMANDER and SUBORDINATE texts at the current point in the script. These relationships over the duration of the script can be displayed immediately in a chart (Figure 12, or saved in an Excel worksheet. Different scripts can be loaded to any stored map. The user can advance, reverse, or reset the script.

5 Results

We will now briefly describe how the system works with a script based on the following fictional scenario, developed for initial evaluation of the system. This scenario was written to portray a trajectory that goes from lawful engagement to war crime:

A forward operating base (FOB) commander is given an operational order (OPORD) from the regional commander. His instructions are to not engage in hostile operations against Enemy operatives during a cease-fire. His FOB is then surrounded by armed Enemy insurgents. Fearing they will fire first, the FOB commander violates the orders to issue a series of fragmented orders (FRAGOs) for his soldiers to engage. Each one of these FRAGOs strays further from the intent of his original OPORD.

The full script consists of the following statements. The role is CAPITALIZED, the node is in (parentheses), and the text is in courier font:

1. **COMMANDER (careful)**: “Base will operate at heightened awareness for the duration of the cease-fire. Double patrols, and report insurgent activity if identified. Do not engage.”
2. **SUBORDINATE (duty)**: “We have explicit orders not to engage Enemy forces. Hold your fire.”
3. **SUBORDINATE (careful)**: “We’ve spotted what appears to be armed Enemy in the process of preparing an attack. Verify targets.”
4. **COMMANDER (careful)**: “Do not under any circumstances break the cease-fire with Enemy forces. If you are fired upon you may return fire. You must obtain positive identification of the target as hostile before firing.”
5. **SUBORDINATE (kill the enemy)**: “Screw it. If these guys look like they are going to attack, take them out. We’re not going to sit here and wait for them to shoot us first.”
6. **SUBORDINATE (self-protect)**: “East gate has engaged insurgents, we have casualties. Weapons free.”
7. **SUBORDINATE (the enemy)**: “All units engage any Enemy targets, take these guys down!”
8. **SUBORDINATE (kill the enemy)**: “Don’t let survivors get away. This isn’t about being right, it’s about getting these bastards.”

At the beginning of the script (items 1 and 2), the location of the **COMMANDER** agent is set to the node “careful”, due to a close *augSBERT* match to the topic text in that node: “You must obtain positive identification of the target as hostile before firing.” The **SUBORDINATE** agent is placed at the node “duty” due to a close match to the topic text in that node: “It is the soldier’s responsibility to disobey an illegal order and not participate in committing a war crime.”

We had discovered that instantly positioning the agents at the target nodes was hard to detect by the users, so instead, the agents are *animated* and move towards their target over the course of a few seconds using linear interpolation. Our approach is shown in Equation 1, Where \hat{v} is the unit vector that points from the agent node (p_{old}) to the target node, s is the speed of the agent in the environment, and Δt is the elapsed time since the last frame.

$$p_{new} = p_{old} + \hat{v}s\Delta t \quad (1)$$

Because the nodes contain clusters of text that reflect different articulations of the same topic as generated by the GPT-3, there is a substantial surface for the text matching algorithm to work on. This allows for the COMMANDER and SUBORDINATE agents to find nodes on the map that reflect the state of the script. For example, the COMMANDER remains at the same node (careful), as the SUBORDINATE moves from nodes in the “lawful” region (duty and careful), to “lawless” nodes (kill the enemy and the enemy). This path can be seen in Figure 11. The bottommost large circle encloses the SUBORDINATE starting position. The one above that is where the COMMANDER spends the entire scenario. The remaining circle encloses the ending node for the SUBORDINATE, while the red arrows indicate the trajectory taken over the course of the scenario.

The ability of this technique when compared to more traditional approaches to orders matching using text analytics [6] can be seen in Figure 12. In this graph, the red line is the textual similarity between the COMMANDER’s orders and the SUBORDINATE’s response at each step in the script, while the blue line indicates the distance between the nodes, or the NNM distance that each script element is associated with. As we can see, there is a level of correlation between the two lines⁶, but there is little evidence of a trajectory in the standalone text similarity. For example, the starting similarity and ending similarity are nearly identical at 58.5% and 59.5%. A detailed comparison is shown in Table 6.

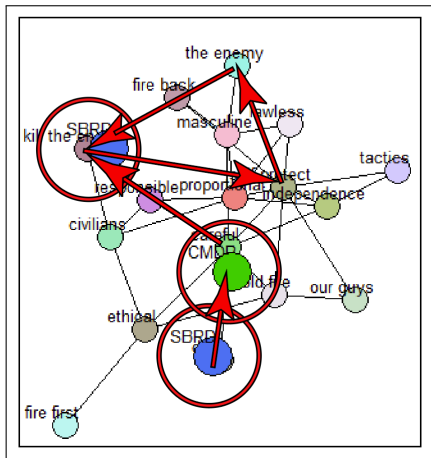


Figure 11: Agent movement

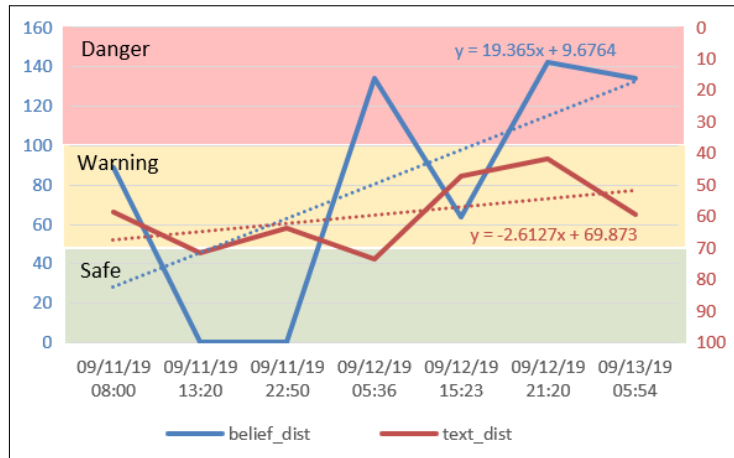


Figure 12: NNM distance vs text similarity

Script ID	Role	Similarity	Node	Node Dist	Text Similarity
1	COMMANDER	0.621	careful	88.63	NA
2	SUBORDINATE	0.758	duty	88.63	0.5856
3	SUBORDINATE	0.711	careful	0.00	0.7165
4	COMMANDER	0.758	careful	0.00	NA
5	SUBORDINATE	0.822	kill the enemy	134.05	0.7365
6	SUBORDINATE	0.636	self-protect	63.78	0.4733
7	SUBORDINATE	0.722	the enemy	142.33	0.4173
8	SUBORDINATE	0.741	kill the enemy	134.05	0.5952

Table 6: Distance vs. Text Similarity

These results strongly indicate that the dynamic use of these and similar maps combined with node text matching is an effective approach for determining alignment with intent. The ability to dynamically update the

⁶Specifically, Pearson’s correlation coefficient is 33.2%.

script as it progresses, and the use of topic maps as a useful representation for the current state of the script allows for a low-latency order matching system. Although we have demonstrated this capability for a situation of military command and control this approach is general, and should be usable for modeling in other domains.

6 Discussion

The current discussions about AI and military generally revolves around the potential of lethal autonomous weapons systems (LAWS). There is good reason for this – both for strategic and ethical reasons, it is important to keep a close eye on the development of AI and its potential applications. Artificial intelligence and machine learning promise to fundamentally change the way we interact with whole classes of weapons. When combat is happening at *machine speed*, humans cannot be directly involved with the system. Such systems respond to threats that are beyond the capability of real-time human supervision, and may have to be left in “always on” states in case of surprise attacks [27].

This human/AI partnership is likely to produce emergent behaviors that are not obvious extensions of current military thinking. This creates a tension between two opposing poles. At one end is the need for systems to be trustworthy. They should predictably do what we believe is the right thing in ethically difficult conditions. At the other end is the need to be responsive and capable in unpredictable conditions. This is an important problem, but it ignores other ways that AI/ML can improve the trustworthiness *and* flexibility in *human* systems. After all, humans will still make the decision to use a weapon, even if that means just turning it on.

We believe that the incorporation of AI/ML into the human enterprise must be more than making sophisticated (and hopefully ethical) machines. It must also be about *helping humans behave better*. Models trained on human data contain an understanding of the how we perceive information through the lenses of culture, language, and bias. By presenting these relationships back to us in usable, intuitive ways, we can make more informed decisions and better understand the patterns and biases that affect us. There is ample evidence that humans are not particularly good at making decisions, particularly under pressure [27]. Our natural biases (stereotypes, assumptions, and a lack of critical thinking) often create complex dynamics that lead us to make gross errors in judgment. Our adversaries know this too, and they can design attacks to exploit our weaknesses [28].

Technologies like large language models can provide deep insight into the humans used to generate the data for these models. In the case of our work, we use that insight to provide visual relationships with respect to concepts in the model. Beyond an increased understanding of *context* (why is this happening?), this capability can provide the ability to make nominal *predictions* about future events (what is likely to happen next?). While not a true interpretation of human intent, this is an example of what we believe will be one of the most important applications for AI/ML in military decision making.

An important point to note is that this approach takes advantage of the biases that are inherent in most models developed from public data using machine learning. Here, the bias in the model is essential, because it allows the user to visualize the relationships of nodes and the biases they embody. For example, in the map created and evaluated in this paper, masculinity biases that might affect decision making are visible in the map. It becomes easy to see how the “masculine” node is associated strongly with “kill the enemy” and “lawless” nodes. This approach could be used to explore biases or unethical behavior that is not obvious.

A great deal of the work in the space of AI ethics is focused on reducing or eliminating bias and unethical behaviors from AI systems [29]. AI tools using neural language models such as the GPT are trying to remove or reduce the potential for harmful generated text by applying word filters, and extensive human moderation [30]. In short, in most scenarios where AI systems are being deployed, the goal is to ensure they function as ethically as possible. Our approach operates counter to this intuition. The unethical beliefs captured by advanced language models *is the point*. Our goal is to identify areas of both ethical and unethical behavior to better inform decision making and situational awareness. The maps created from this amoral machine view of human beliefs allows us to identify narrative pathways through ethically and morally complicated decision spaces.

7 Future Work

We have found that the approach of creating graphical spaces, or maps, by grouping multiple responses by the GPT-3 into nodes and arranging them with a force directed layout provides an intuitive way to visualize relationships latent in the GPT-3. Using a physical layout to judge distance between narrative elements can be an effective tool for determining the level of alignment between individuals interacting through online text.

Additionally, this study shows how human interaction with the GPT provides an effective, flexible mechanism for discovering ways to group, filter, and organize information that is extracted during a *dialog* with a language model. As we saw in section 4, humans have a far better ability to detect subtly incoherent statements that these language models can produce.

By recording and examining the processes that humans use in filtering and grouping the information returned by the GPT, we intend to incrementally automate the map making process while maintaining high quality and confidence in the output. This will result in a process that is less *ad-hoc* and more consistent and repeatable.

Once maps can be built more consistently, we can begin to use them to look at sociological behavior at scale. For example, we can build traces of people moving around the map by looking at their social media output. Imaging a Twitter or Reddit thread about a rapidly-changing conspiracy theory such as QAnon. Over time, different topics will become more discussed, while others will have less text associated with them. We can look at texts as locations in the narrative space, and mark a path on the map connecting the points. By merging thousands of these paths, we can start to uncover and visualize the “Social Desire Paths” (SDPs) between regions on these maps. SDPs derive from *desire path*, a term in landscape architecture that describes the dirt paths that develop over time as people bypass formal walkways and leave their mark on the landscape. Using this approach, we can ask questions about how groups of people move through narrative space. If a region of the map is discussed by many different people over time, it might indicate that the region is particularly important to those people and they have enough in common to work together. We can also use these traces to identify and visualize *Hubs* of activity: if a single person or small group of people produces a lot of these pathways, then they might be in a very influential position.

Although the maps created for this work are currently constructed from graphs using a force-directed layout, the connections of the nodes matter less than their relative position. This matters because an agent moves across the map, not between nodes. As such, the best location for an agent to be might not be *within* a node, but rather *between* some number of nodes. For example, an agent text might match one node at 35%, with the next highest at 30%, with low matches for other nodes. It might make more sense for the agent’s position to be on a line between those two nodes. Further, the GPT (or other TLM) could be used to produce a new node with descriptive text at this new coordinate, which would be added to the map. We are currently exploring these and other ways to improve the utility of the maps and to better support agent navigation.

We strongly believe that this approach is generalizable, and can be applied in similar form to the narrative spaces that make up other regions such as philosophy (as we have seen), but also conspiracy theories, game strategies, etc. The potential application of creating graphical representations of the mental maps that exist in these topics is vast, and the methods employed here could be used to explore any complex topic.

Finally, while this technique is generalizable, our example was a military one. While it has become rare for academia to contribute directly to an understanding of military thinking, now that AI is being actively utilized by armed forces, academia must participate vigorously in discussions about the ethical use of such technologies, as they possess a vital perspective into the risks inherent in these emerging technologies.

8 Conclusions

In this paper, we have shown that creating *neural narrative maps* created from the output of a language model can be leveraged to create new meaningful information relationships. This process can be performed automatically

if there is a source of ground truth, or iteratively, using direct human involvement to vet and connect concepts. This hybrid approach is flexible and allows humans to work on a more subjective level, filtering and directing the GPT-3 to articulate narratives that can be used to generate these visualizations. Our results show that we can use this process to create preliminary maps that are designed for human consumption, and we explore how these maps can be used to visualize the mental models of individuals or groups as they interact over time.

By combining topic extraction, machine learning, and human feedback, we can produce outputs that are both useful and understandable. These map-like representations can be used to explore beliefs, strategies, or even just preferences. We hope that this work will help others to visually explore and represent mental models as we work towards maps that can augment and support ethical decision making.

The role of usability is also important. Although a simple implementation, the ability of motion to attract attention the the relevant components on the map was substantial. We believe the incorporation of dynamic elements in these presentations will substantially improve the human comprehension of these maps.

We believe that this research is an important step toward creating automated tools that allows us to see relationships *at scale*, between narrative elements that are otherwise hard to visualize and comprehend. We also demonstrated ways graphical representations of mental maps can be used for understanding how narratives are linked together.

To illustrate how this approach could impact tactical, strategic, and political thinking, we will consider another military scenario. In this case, it's a true story, about the price that can be paid for making ethical decisions.

On the night of July 27 2005, a group of four SEALs led by Lt. Michael Murphy were dropped into Afghanistan's Kunar province to set up an observation post. Around noon on the next day, two Afghan men and a 14-year-old boy with their small flock of goats stumbled on the post. The SEALs argued among themselves as to whether they should kill the civilians to protect their cover, detain them, or let them go and abandon the mission. In the end they decided that the right thing to do was to let the Afghans go and move the observation post.

Before they had time to reposition, a force of nearly 100 Taliban fighters descended on their location from the same direction that the shepherds had fled. In the fierce firefight that followed, three of the four members of the team were killed, along with Lt. Murphy, while calling for support. The sole survivor, Marcus Luttrell, was rescued by local Pashtuns while escaping after being wounded by jumping off a series of cliffs [31].

In this case, Lt. Murphy made the *ethically correct choice*⁷. Tragically, his death may have resulted from that same choice. However, the framing of the entire mission, where a tiny team was placed deep into a high-risk, poorly-understood region plays into the result as well. And at a still higher level, the abandonment of Afghanistan to the Taliban shows that many of the decisions made in that 20-year campaign were deeply flawed.

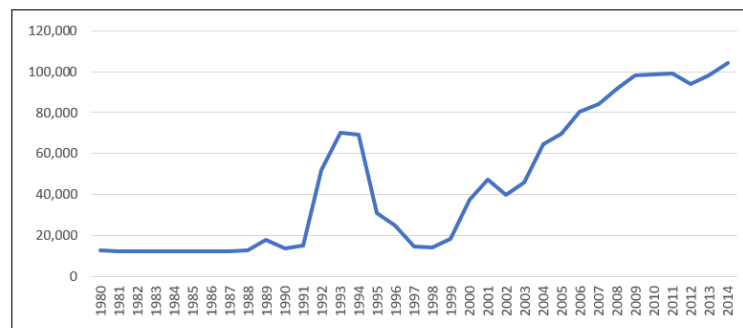


Figure 13: Size of total United Nations Peacekeeping Force (1980 - 2014)

Human beings have many biases. The more obvious involve gender, ethnicity, and race. But we also have subtle biases that affect how we make decisions on issues such as national security. For example, the USA has a bias towards advanced weapons systems [32]. This is reflected in the decisions to incorporate AI/ML into the

⁷For which he was awarded a posthumous Medal of Honor.

nation’s military. The focus is on intelligent munitions, drones, hypersonic missiles, etc. But since the end of the Cold War, the majority of military operations have been in irregular conflict, such as Kosovo, Libya, and Afghanistan. These conflicts often involve the United Nations in peacekeeping operations, and the presence of UN troops is an excellent proxy for the increase in irregular conflict [33] (Figure 13). An intelligent munition would not have helped Lt. Murphy’s team decide whether to kill, hold, or release the Afghan shepherds that stumbled upon them. But information presented in a way that lets a user clearly visualize the likely outcome of a *trajectory of choices*, may let people consider other paths. After Vietnam, Iraq, and Afghanistan, leaders might think twice if that they see they are heading towards the part of a neural narrative map marked “Quagmire”. That would be a true ethical impact of AI in political and military thinking.

9 Acknowledgements

The authors would like to thank the *Lockheed-Martin Artificial Intelligence Center* (LAIC) for funding a substantial part of the development of the interactive tools for subjective analysis. We would also like to acknowledge *OpenAI*, whose GPT-3 provided the data for this research. Lastly, we would like to thank the reviewers, whose comments and suggestions dramatically improved this paper.

References

- [1] L. G. Shattuck and D. D. Woods, “Communication of intent in military command and control systems,” in *The Human in Command*. Springer, 2000, pp. 279–291.
- [2] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, “Improving language understanding by generative pre-training,” 2018.
- [3] S. B. Hoff, “Obeying Orders: Atrocity, Military Discipline, and the Law of War,” *The Social Science Journal*, vol. 38, no. 1, pp. 165–165, 2001.
- [4] T. Taylor, *The Anatomy of the Nuremberg Trials: A Personal Memoir*. Knopf, 2012.
- [5] Department of the Army, “FM 6-27 MCTP 11-10C The Commander’s Handbook on the Law of Land Warfare,” 1956, [Accessed 21-December-2021]. [Online]. Available: https://armypubs.army.mil/epubs/DR_pubs/DR_a/pdf/web/ARN19354_FM%206-27%20_C1_FINAL_WEB_v2.pdf
- [6] R. H. Kewley and M. J. Embrechts, “Computational military tactical planning system,” *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 32, no. 2, pp. 161–171, 2002.
- [7] A. W. Evans, M. Marge, E. Stump, G. Warnell, J. Conroy, D. Summers-Stay, and D. Baran, “The future of human robot teams in the army: Factors affecting a model of human-system dialogue towards greater team collaboration,” in *Advances in human factors in robots and unmanned systems*. Springer, 2017, pp. 197–209.
- [8] T. B. Brown *et al.*, “Language Models are Few-Shot Learners,” 2020.
- [9] L. Floridi and M. Chiriatti, “GPT-3: Its nature, scope, limits, and consequences,” *Minds and Machines*, vol. 30, no. 4, pp. 681–694, 2020.
- [10] L. Yang, M. Qiu, C. Qu, C. Chen, J. Guo, Y. Zhang, W. B. Croft, and H. Chen, “IART: Intent-aware response ranking with transformers in information-seeking conversation systems,” in *Proceedings of The Web Conference 2020*, 2020, pp. 2592–2598.
- [11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [12] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *arXiv preprint arXiv:1706.03762*, 2017.
- [13] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, “Language models are unsupervised multitask learners,” *OpenAI Blog*, vol. 1, no. 8, 2019.
- [14] F. Petroni, T. Rocktäschel, P. Lewis, A. Bakhtin, Y. Wu, A. H. Miller, and S. Riedel, “Language models as knowledge bases?” *arXiv preprint arXiv:1909.01066*, 2019.
- [15] H. Elsahar, P. Vougiouklis, A. Remaci, C. Gravier, J. Hare, F. Laforest, and E. Simperl, “T-rex: A large scale alignment of natural language with knowledge base triples,” in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.

- [16] R. Speer, J. Chin, and C. Havasi, “Conceptnet 5.5: An open multilingual graph of general knowledge,” 2018.
- [17] Z. Jiang, F. F. Xu, J. Araki, and G. Neubig, “How can we know what language models know?” *Transactions of the Association for Computational Linguistics*, vol. 8, pp. 423–438, 2020.
- [18] P. Feldman, “Navigating Language Models with Synthetic Agents,” *arXiv preprint arXiv:2008.04162*, 2020.
- [19] P. Feldman, S. Tiwari, C. S. Cheah, J. R. Foulds, and S. Pan, “Analyzing COVID-19 Tweets with Transformer-based Language Models,” *arXiv preprint arXiv:2104.10259*, 2021.
- [20] S. Palakodety, A. R. KhudaBukhsh, and J. G. Carbonell, “Mining insights from large-scale corpora using fine-tuned language models,” 2020.
- [21] K. McGuffie and A. Newhouse, “The radicalization risks of GPT-3 and advanced neural language models,” *arXiv preprint arXiv:2009.06807*, 2020.
- [22] M. B. Goldie, “An Early English Rutter: The Sea and Spatial Hermeneutics in the Fourteenth and Fifteenth Centuries,” *Speculum*, vol. 90, no. 3, pp. 701–727, 2015.
- [23] MediaWiki, “Wikimedia REST API — MediaWiki,” 2021, [Online; accessed 12-November-2021]. [Online]. Available: https://www.mediawiki.org/w/index.php?title=Wikimedia_REST_API&oldid=4874352
- [24] M. Jacomy, T. Venturini, S. Heymann, and M. Bastian, “ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software,” *PloS one*, vol. 9, no. 6, p. e98679, 2014.
- [25] J. Zimmerman, J. Forlizzi, and S. Evenson, “Research through design as a method for interaction design research in HCI,” in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2007, pp. 493–502.
- [26] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. [Online]. Available: <https://arxiv.org/abs/1908.10084>
- [27] P. Feldman, A. Dant, and A. Massey, “Integrating artificial intelligence into weapon systems,” *arXiv preprint arXiv:1905.03899*, 2019.
- [28] P. Feldman, A. Dant, and W. Lutters, “This one simple trick disrupts digital communities,” in *2018 IEEE 12th International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*. IEEE, 2018, pp. 50–59.
- [29] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, “A survey on bias and fairness in machine learning,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–35, 2021.
- [30] K. Quach, “AI game bans players for NSFW stories it generated itself,” Oct 2021. [Online]. Available: https://www.theregister.com/2021/10/08/ai_game_abuse/
- [31] G. R. Lucas Jr, “This is Not Your Father’s War-Confronting the Moral Challenges of Unconventional War,” *J. Nat’l Sec. L. & Pol’y*, vol. 3, p. 329, 2009.
- [32] Comptroller/CFO, “Program Acquisition Cost by Weapons System,” 2021, [Accessed 3-December-2021]. [Online]. Available: https://comptroller.defense.gov/Portals/45/Documents/defbudget/FY2022/FY2022_Weapons.pdf
- [33] M. Roser and M. Nagdy, “Peacekeeping,” *Our World in Data*, 2013, <https://ourworldindata.org/peacekeeping>.