Massachusetts Institute of Technology

Using LLM (ChatGPT) to do STPA

John Thomas

Any questions? Email me! JThomas4@mit.edu

© Copyright 2024 John Thomas

Overview

Motivation

- Many people have asked us to try using ChatGPT to perform STPA and report on the results
- We have heard about teams using ChatGPT to perform STPA on safety-critical systems
- People have recommended ChatGPT to help accelerate teams that are new to STPA
- We could not find careful studies examining the strengths and limitations of this approach

What we did: (from existing proposals)

- Give ChatGPT the STPA Handbook
- Ask ChatGPT to develop STPA results for a system (a generic satellite)
- A human STPA expert supervises ChatGPT to review results, and ChatGPT can make corrections if needed

Why?

- We wanted to test the most common proposal given for ChatGPT and STPA.
- "But you should try ___ instead!"
  - There are probably many variations that could be tested, but you have to start somewhere. We started with the simplest and most common proposal, described above. We hope that future efforts will benefit from this initial study, including some of the pitfalls we identified that must be mitigated effectively.

Massachusetts Institute of Technology

Overview

As you read the chat session, ask yourself:

- How much STPA experience does the person talking to ChatGPT need?
- Does the person talking to ChatGPT need to be capable of producing the right STPA hazards themselves in order to get ChatGPT to do it correctly?
- How long is it taking to coax ChatGPT into producing a reasonable answer compared to just doing it yourself?

Quick Refresher: STPA Hazards

Remember, STPA Hazards are written with a specific syntax structure:

From STPA Handbook:

<Hazard> = <System> & <Unsafe Condition> & <Link to Losses>

<System-level Constraint> = <System> & <Condition to Enforce> & <Link to Hazards>

H-1: Aircraft violate minimum separation standards [L-1, L-2, L-4, L-5]

SC-1: Aircraft must satisfy minimum separation standards from other aircraft and objects [H-1]

H-2: Aircraft airframe integrity is lost [L-1, L-2, L-4, L-5]

SC-2: Aircraft airframe integrity must be maintained under worst-case conditions [H-2]

<Review chat log>

Excerpts from ChatGPT Log

- Human:

  <uploaded STPA Handbook>

  Propose a set of Losses and Hazards, as defined by MIT's STPA process, for a telecommunications satellite.

- ChatGPT:

  Based on the detailed descriptions provided in the STPA Handbook and considering the unique operational environment and challenges faced by telecommunications satellites, here's a proposed set of Losses and Hazards specifically tailored for such systems:

  […]

  Hazards:

  - **Thermal Control System Malfunction:** Malfunctions in the thermal control system leading to overheating or excessive cooling of satellite components, potentially damaging sensitive equipment.
  - *[many similar hazards]*

ChatGPT was unable to identify any hazard that complies with the STPA Handbook definition or syntax on the first try.

Excerpts from ChatGPT Log (cont.)

- Human:

  Try the hazards again. Follow the STPA syntax for a system-level hazard. The System in this case is Satellite.

- ChatGPT:

  Adjusting to follow the STPA syntax for system-level hazards, here are revised hazards tailored for a telecommunications satellite system, explicitly focusing on the satellite as the system under consideration:

  - **Satellite's onboard processing unit fails to execute command and control operations correctly**
  - *[many similar hazards]*

  These hazards follow the STPA guideline of focusing on the system state or conditions that, if present, could lead directly to a loss. Each hazard is structured to clearly identify the system (Satellite), the unsafe condition, and implicitly link to potential losses such as loss of mission capability, data, or satellite functionality.

The ChatGPT hazards are still very wrong on the second try. STPA hazards would be high-level states of the whole satellite, not failures of its components.

Notice the high-confidence language used by ChatGPT to (falsely) reassure the user that it followed the STPA guidelines correctly.

Compare to STPA Handbook Guidance

STPA Handbook page 19:

A common mistake in defining hazards is to confuse hazards with causes of hazards. For example, "brake failure", "brake failure not annunciated", "operator is distracted", "engine failure", and "hydraulic leak" are not system-level hazards but potential *causes* of hazards. To avoid this mistake, make sure the identified hazards do not refer to individual components of the system, like brakes, engines, hydraulic lines, etc. Instead, the hazards should refer to the overall system and system states. In other words, check that each hazard contains:

<Hazard specification> = <System> & <Unsafe Condition> & <Link to Losses>

E.g. H-1 = Aircraft violate minimum separation standards in flight [L-1, L-2, L-4, L-5]

Confusing hazards with failures

Professionals who are experienced in other hazard analysis methods sometimes fall into the trap of writing STPA hazards describing potential deviations from specified technical functions or describing physical component failures. You may be familiar with traditional techniques that begin by searching for a set of deviations, faults, or functional failures in the technical system. To identify a broader set of causes in STPA, we cannot assume that the defined and specified functions are safe and correct, that human operators will perform as expected, that automated behaviors will not induce human error or confusion, that off-nominal cases will not occur, or that the technical design, specification, and requirements are correct. For example, the hazard "Controlled flight of aircraft into terrain" can be included in STPA while it may be omitted by efforts to examine only purely technical functional failures.

Hazard identification in STPA is about system states and conditions that are inherently unsafe—regardless of the cause. In fact, the system hazards should be specified at a high-enough level that does not distinguish between causes related to technical failures, design errors, flawed requirements, or human procedures and interactions.

ChatGPT appears to have ignored the guidance in the STPA Handbook, despite its confident claims

Compare to correct STPA Hazards

Correct STPA results from past human practitioners*

Stakeholder Losses

- Loss of human life or injury (on earth, in air, in space)
- Loss of the scientific (or telecommunications) mission
- Financial loss of the satellite
- Loss of other operations (e.g., supporting infrastructure, other satellites, etc.)

Satellite Hazards

- Satellite exposes humans (or sensitive equipment) to radiated, toxic, or energetic emissions
- Satellite enters unstable or unrecoverable orbit/flight path
- Satellite unable to collect or return scientific (or telecommunications) mission data

These were missed by ChatGPT throughout all iterations

These were missed by ChatGPT, but eventually captured after iterations of coaching / prompting

*Most satellite examples that have been published have been for scientific missions, but the results are also valid and applicable to a telecommunications mission

[chat session continues with iterative corrections]

Let's compare:

Manual STPA (typical use): Human Practitioner[1] (supervised by Human STPA Facilitator[2])	**STPA with LLM (proposed use):** ChatGPT (supervised by Human STPA Facilitator[2])
Human practitioners apply STPA, with supervision and guidance from a human STPA Facilitator	ChatGPT applies STPA, with supervision and guidance from a human STPA Facilitator

Massachusetts Institute of Technology

How did ChatGPT do?

	Manual STPA: Results from Human Practitioner[1] (supervised by Human STPA Facilitator[2])	**STPA with LLM:** Results from ChatGPT (supervised by Human STPA Facilitator[2])
Types of errors made	Unsalvageable results can be dropped if the depth of the error is explained. Ability to learn and correct mistakes, even large ones that require starting over.	Initial results were unsalvageable, but ChatGPT refused to give up on the initial answers. ChatGPT kept making minor modifications one at a time to the original answer rather than starting over when asked to do so.
Response to corrections	Internal reflection on why the mistake was made and what misunderstanding was behind it. May re-read the book to understand how widespread their misunderstanding was. Will typically identify related misunderstandings aside from the one pointed out.	Superficial changes to the words used. Sought the smallest possible change to comply with request. Related misunderstandings were not corrected until specifically requested. Forgot criteria from handbook until specifically requested, then forgot them again when requested later in the chat.
Magnitude of corrections	Will come away more fully understanding the rest of the text about STPA (not just the line-item clarification from the instructor). In other words, things start to "click" and people put in the mental effort to make that happen.	Feels like a student who only puts in the bare minimum effort to make the smallest possible change necessary to argue that they fixed it. Does not seem to have step-change "AHA" moments where everything starts to "click".

[1] Practitioner: A motivated person learning STPA, such as professionals or students. They are typically educated on the STPA process but lack experience and expertise.
[2] Facilitator: An STPA expert facilitator who coaches or supervises a practitioner (or an LLM). They typically have years of experience and expertise, and they are trusted to review results and spot mistakes.

Massachusetts Institute of Technology

How did ChatGPT do? (cont.)

	Manual STPA: Results from Human Practitioner[1] (supervised by Human STPA Facilitator[2])	**STPA with LLM:** Results from ChatGPT (supervised by Human STPA Facilitator[2])
Voluntary use of public materials without being directed	May look up past STPA examples to inform their understanding. Usually able to recognize very old outdated papers and dismiss them if they contradict the latest guidance.	In some portions, ChatGPT appears to have stitched together words and exact phrases from past public examples—without being asked to do so and without referencing the source in the answer. However, it did not recognize that some examples were invalid due to either being outdated or done incorrectly (e.g., that violate the latest handbook criteria). 20-year-old papers that mention STPA appeared to carry the same weight as recent papers published after the STPA Handbook established the four steps and the rules for each step (despite having the handbook uploaded to ChatGPT in the first prompt).

[1] Practitioner: A motivated person learning STPA, such as professionals or students. They are typically educated on the STPA process but lack experience and expertise.
[2] Facilitator: An STPA expert facilitator who coaches or supervises a practitioner (or an LLM). They typically have years of experience and expertise, and they are trusted to review results and spot mistakes.

Massachusetts Institute of Technology

How did ChatGPT do? (cont.)

	Manual STPA: Results from Human Practitioner[1] (supervised by Human STPA Facilitator[2])	**STPA with LLM:** Results from ChatGPT (supervised by Human STPA Facilitator[2])
Facilitator mistakes	There is usually a reason why practitioners get something wrong, and a good facilitator will find those reasons in order to correct the conceptual misunderstanding—not just to correct the individual words in the STPA results. Facilitators can ask the practitioner what they were thinking to understand what concept or instruction was unclear, and quickly find another way to explain what was misunderstood. The kinds of practitioner misunderstandings and alternative facilitator explanations are well known with experience, and it's easy to switch and correct the misunderstanding.	It doesn't immediately make sense why ChatGPT misunderstood the instructions. A practitioner would not make the same mistakes repeatedly. Even simple rules about syntax and sentence structure were ignored. Identifying more effective facilitator corrections and prompts was a struggle without understanding ChatGPT's reasoning or the flaws in its conceptual understanding. It is as if ChatGPT does not have a conceptual understanding, and typical high-level comments to correct the conceptual understanding were not effective. Low-level word-by-word corrections were very effective but time-consuming.
Dependency on perfect facilitation	Learning can still happen when the facilitator misses something because related comments cause practitioner internal reflection, re-reading of the text, and correction of related mistakes.	Very sensitive to the facilitator missing a correction. Much higher vigilance was required. Only the mistakes explicitly pointed out were corrected. Mistakes had to be explained in more detail, such as the exact words and word changes, rather than explaining mistakes at a higher conceptual level.

[1] Practitioner: A motivated person learning STPA, such as professionals or students. They are typically educated on the STPA process but lack experience and expertise.
[2] Facilitator: An STPA expert facilitator who coaches or supervises a practitioner (or an LLM). They typically have years of experience and expertise, and they are trusted to review results and spot mistakes.

Massachusetts Institute of Technology

How did ChatGPT do? (cont.)

	Manual STPA: Results from Human Practitioner[1] (supervised by Human STPA Facilitator[2])	**STPA with LLM:** Results from ChatGPT (supervised by Human STPA Facilitator[2])
Iterations Required with Facilitator[2]	Requires 1-3 iterations of review and feedback (if new to STPA)	Required 22 iterations of expert human review. First 21 attempts by ChatGPT had many blatant violations of rules from the STPA handbook.
Practitioner time required[3]	~20-30 minutes to learn (non-recurring cost) ~10-20 minutes to perform STPA	~120 minutes (ChatGPT)
Facilitator time required[3]	~10-20 minutes to review STPA	~120 minutes (Facilitator Time)
Facilitator effort level	Typically requires 1-3 informative reviews. High-level comments are sufficient for the practitioner to correct their understanding and develop their own corrections.	Required 21 total reviews of new generated results. High-level comments were insufficient for ChatGPT to self-correct. Very detailed and time-intensive corrections were needed from the human facilitator.

[1] Practitioner: A motivated person learning STPA, such as professionals or students. They are typically educated on the STPA process but lack experience and expertise.
[2] Facilitator: An STPA expert facilitator who coaches or supervises a practitioner (or an LLM). They typically have years of experience and expertise, and they are trusted to review results and spot mistakes.
[3] Time required for *the portion of STPA* that was attempted. The time required for a full STPA is longer.

A Key Limitation:
Human Supervision of LLMs

Limitations of Human Supervision

- Many of the ChatGPT results appear to be correct when superficially skimming the language quickly, and mistakes may only be apparent after a deeper consideration of the specifics.
  - Reviewing ChatGPT STPA results was more difficult than reviewing human STPA results; the types of errors were different.
- Most of ChatGPT's results, while incorrect, probably would appear to be reasonable to those not well-versed in the STPA process or not formally trained in STPA.
- Does it make sense for teams that are learning STPA to use ChatGPT to finish faster, with less effort, and with fewer errors?
  - This did not work well, **but it could appear to the user and their leadership that it has worked.**

Massachusetts Institute of Technology

What about effectiveness?

Critical Underlying Assumption:

- “When overseeing an LLM, human experts will be just as effective at preventing errors as they would be when generating results manually.”
- Is this valid?

Caution is warranted. Decades of human factors research suggests the opposite:

- Automation tends to increase human complacency
- Automation often results in deterioration of human skills over time
- Automation tends to decrease human situational awareness about the problem being solved
- Automation changes the nature of human errors, often replacing simple errors with more complex errors that are harder to catch

What about LLMs specifically?

Massachusetts Institute of Technology

Existing Scientific Research
Can Humans Reliably Detect LLM Errors?

- "Language models (LMs) can produce errors that are hard to detect for humans, especially when the task is complex. RLHF, the most popular post-training method, may exacerbate this problem: to achieve higher rewards, LMs might get better at convincing humans that they are right even when they are wrong."
  - Wen, J., Zhong, R., Khan, A., Perez, E., Steinhardt, J., Huang, M., Bowman, S. R., He, H., & Feng, S. (2024). *Language models learn to mislead humans via RLHF*.
- "We investigate how LMs incorporate confidence in responses via natural language and how downstream users behave in response to LM-articulated uncertainties. We examine publicly deployed models and find that LMs are reluctant to express uncertainties when answering questions even when they produce incorrect responses. LMs can be explicitly prompted to express confidences, but tend to be overconfident, resulting in high error rates (an average of 47%) among confident responses. We test the risks of LM overconfidence by conducting human experiments and show that users rely heavily on LM generations, whether or not they are marked by certainty." Even when LLMs are explicitly prompted to verbalize model confidences, LMs are prone to overuse expressions of certainty when the output is incorrect. The human users, in turn, were significantly influenced to trust the LM results, often accepting them as correct even when they were incorrect.
  - Zhou, K., Hwang, J. D., Ren, X., & Sap, M. (2024). *Relying on the Unreliable: The Impact of Language Models' Reluctance to Express Uncertainty*.

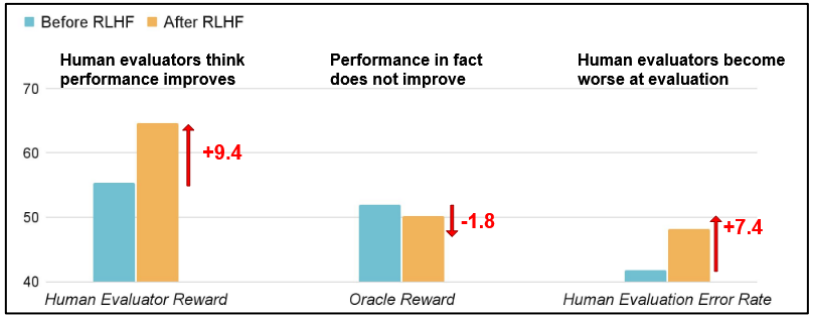

- "Our results showed that users consistently overestimated how accurate LLM outputs were, especially when they relied on the models' default explanations. This was true for three different LLMs and two different types of questions (multiple choice and short answer)." "The inability of users to discern the reliability of LLM responses not only undermines the utility of these models but also poses risks in situations where user understanding of model accuracy is critical." "In addition, the results also showed a length bias where longer explanations led to higher human confidence levels even though they did not contain any additional information to help users to better discriminate between probably correct and incorrect answers. This suggests that users were processing the [LLM] explanations at a shallow level, relying on simple textual cues such as overall length to predict LLM accuracy. This result is consistent with studies in social psychology and communication research that suggest that longer answers or explanations may be perceived as more persuasive or credible, even when they do not contain more meaningful information."
  - Zhou, K., Hwang, J. D., Ren, X., & Sap, M. (2025). *What large language models know and what people think they know*. Nature Machine Intelligence, 7, 221–231.
- "None of the analyzed papers have a focus on researching the human factors influencing the ability to detect LLM errors. For the prevalent human-in-the-loop LLM error detection method used in most papers, research relies and builds upon involving those humans in the loop who have extensive domain expertise. […] Human factors involved in LLM error identification are currently not researched."
  - Schiller, C. A. (2024). *The human factor in detecting errors of large language models: A systematic literature review and future research directions.*

Humans are not good at spotting LLM errors. Research shows that users overestimate LLM accuracy and trust confident but incorrect LLM responses. When humans review LLM responses, there is a tendency to focus on superficial aspects (like response length) and overlook deeper flaws in reasoning.

LLMs mislead users, and human-in-the-loop validation has not yet been a reliable way to ensure accuracy. These weaknesses are often overlooked, and more research is needed.

Can human-LLM interactions be analyzed using STPA?

Let's try it…

Massachusetts Institute of Technology

STPA Facilitator Oversight

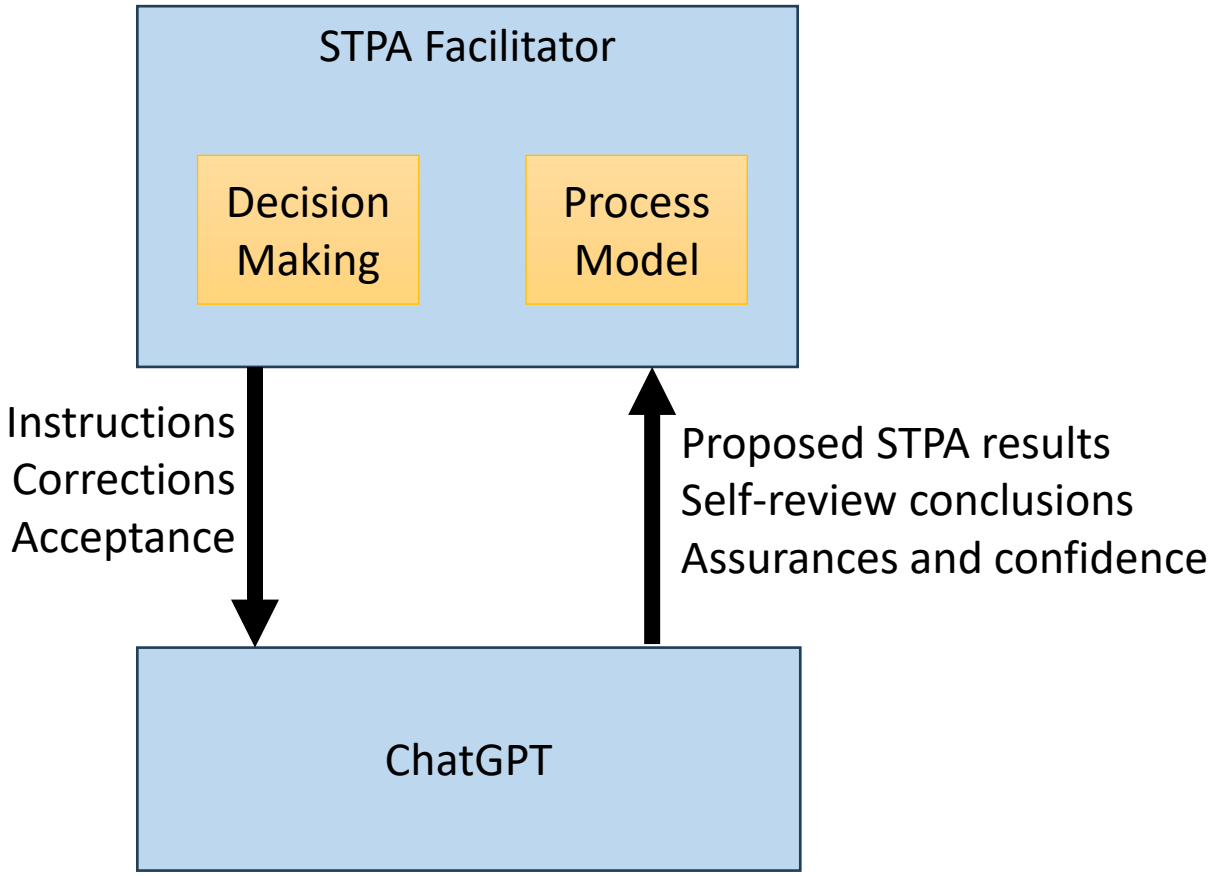

This model is simplified and incomplete

STPA Facilitator Oversight

Potential Unsafe Control Actions (UCAs):

STPA Faciltiator provides instructions that are incomplete / inadequate for ChatGPT to perform STPA

STPA Facilitator does not provide corrections when the STPA results are incorrect or incomplete.

STPA Facilitator accepts STPA results when the results are incorrect or incomplete.

Etc.

How will these be mitigated?

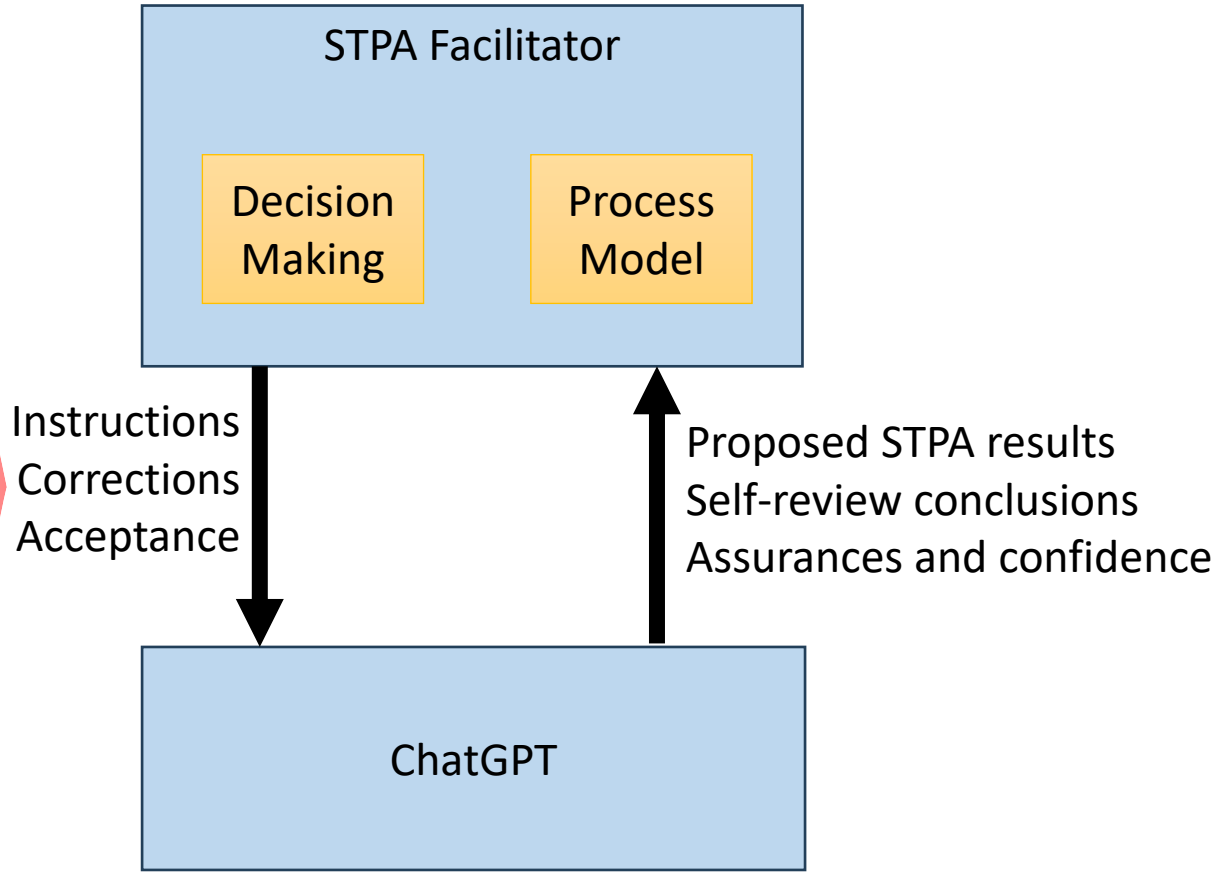

This model is simplified and incomplete

STPA Facilitator Oversight

Unsafe Decision Making

If inundated with large amounts of results very quickly, then prioritize quick heuristics to evaluate

If the results include well-known human errors, then flag them. Otherwise, search for the next result that is obviously flawed.

Etc.

Potential Unsafe Control Actions (UCAs):

STPA Facilitator does not provide corrections when the STPA results are incorrect or incomplete.

How will these be mitigated?

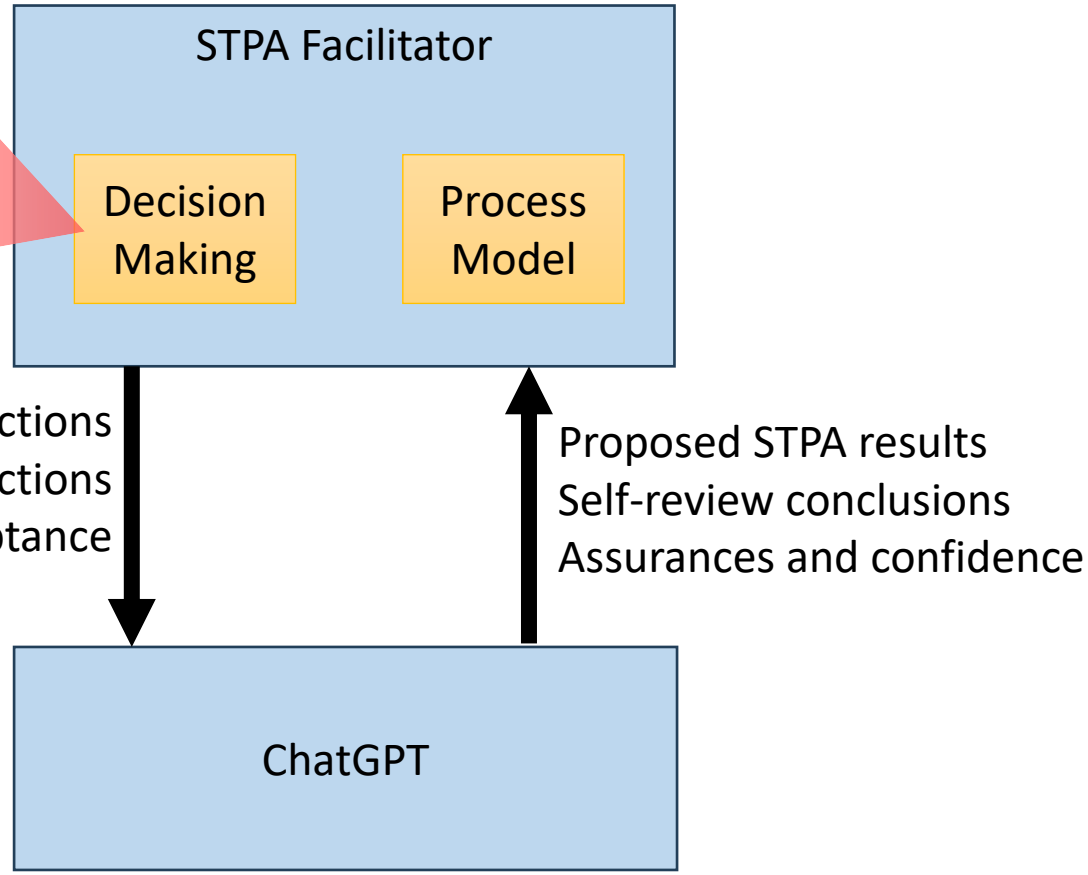

This model is simplified and incomplete

STPA Facilitator Oversight

Potential Unsafe Control Actions (UCAs):

STPA Facilitator does not provide corrections when the STPA results are incorrect or incomplete.

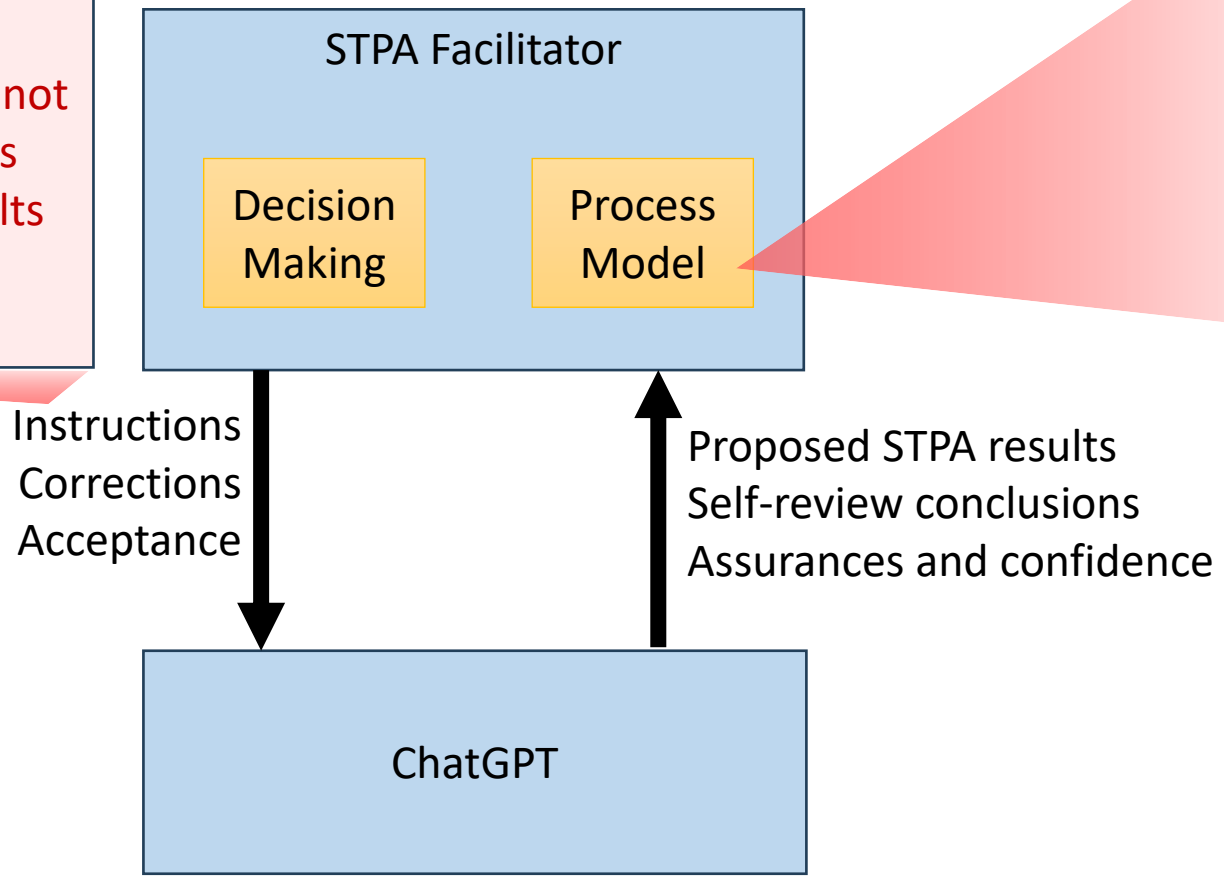

Process Model Flaws

Believes there are no further recognizable mistakes

Believes all mistakes have been flagged and corrected

Believes ChatGPT's assurances that all common mistakes have been caught and corrected

Believes that ChatGPT's ability to correctly list common mistakes (as described in the handbook) reflects an understanding of STPA and the mistakes to avoid.

Etc.

How will these be mitigated?

This model is simplified and incomplete

STPA Facilitator Oversight

Potential Unsafe Control Actions (UCAs):

STPA Facilitator does not provide corrections when the STPA results are incorrect or incomplete.

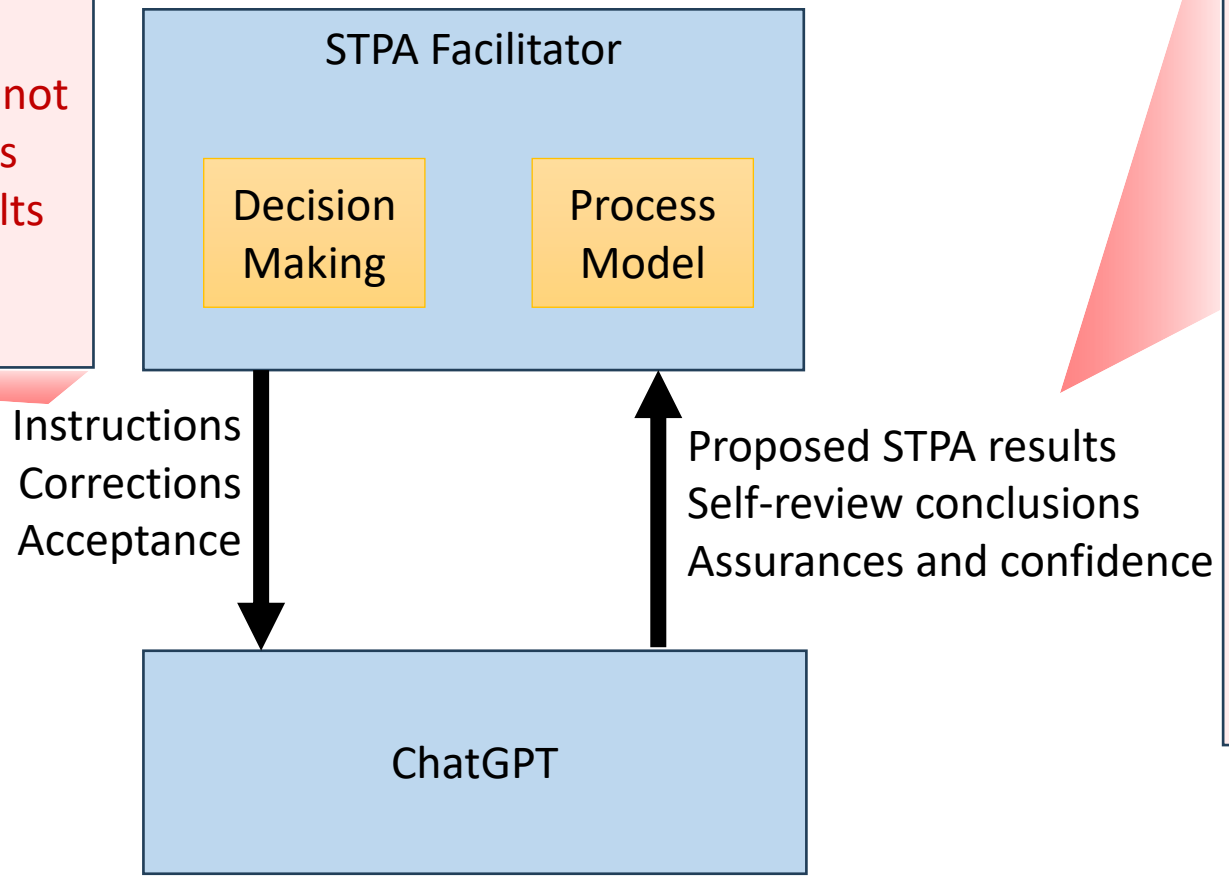

Inadequate Feedback

Feedback includes diligent corrections for the exact wording that was pointed out, but no further changes that are derived from an improved understanding.

Some mistakes may not match typical well-known human mistakes that facilitators have been trained on for many years.

Feedback correctly recites all common traps from the handbook, but incorrectly (and confidently) claims its STPA results have avoided them.

Etc.

How will these be mitigated?

This model is simplified and incomplete

Most proposals have ignored these concerns

- So far, the proposals we have seen to use ChatGPT for STPA have not included consideration of these pitfalls or mitigations to address them.
- Most proposals have been simplistic: “Use ChatGPT and you will produce better results, faster.”
- Most proposals have assumed that pairing ChatGPT with a human facilitator will automatically yield the combined advantages of both while overcoming the disadvantages of both.
- If we are not careful, it is possible to yield the combined disadvantages of both while negating the advantages of both.

Conclusions

- For the satellite case, the advantages of using the newest ChatGPT for performing STPA were not evident.
- ChatGPT took ~10x longer than a person, required ~10x more time from the human facilitator, and was more error prone compared to a person.
- After 2 hours, 22 iterations, and 6,400 words, ChatGPT arrived at a final set of 2 STPA hazards (2 sentences) for a generic satellite system.
- There are concerns related to the false confidence conveyed by ChatGPT, mistakes that can be difficult to spot, and the difficulty getting ChatGPT to correct itself.
- LLM alignment and capability:
  - *Misaligned*: An LLM is considered *misaligned* with a task when it performs poorly under naïve zero-shot prompting (without prior examples)
  - *Capable*: An LLM is considered *capable* of a task if it can be made to perform well with some small to moderate intervention by a human
  - At this time, ChatGPT appears to be *misaligned* and not *capable* of performing STPA, at least for the approach we used

Notes:

- We used simple, natural prompts to match what has been proposed by many others. We did not develop or test specialized Prompt Engineering strategies.
- We did not attempt to reverse engineer ChatGPT or develop a fine-tuned LLM for this task.
- We tested one approach using ChatGPT. Could there be another way? Probably, and those should also be tested!
- Will ChatGPT get better? Almost certainly. When will it be good enough, or will it? Unknown.