


Verkada 0524 AI Powered Search Camera



Verkada 0524 AI Powered Search Camera User Guide

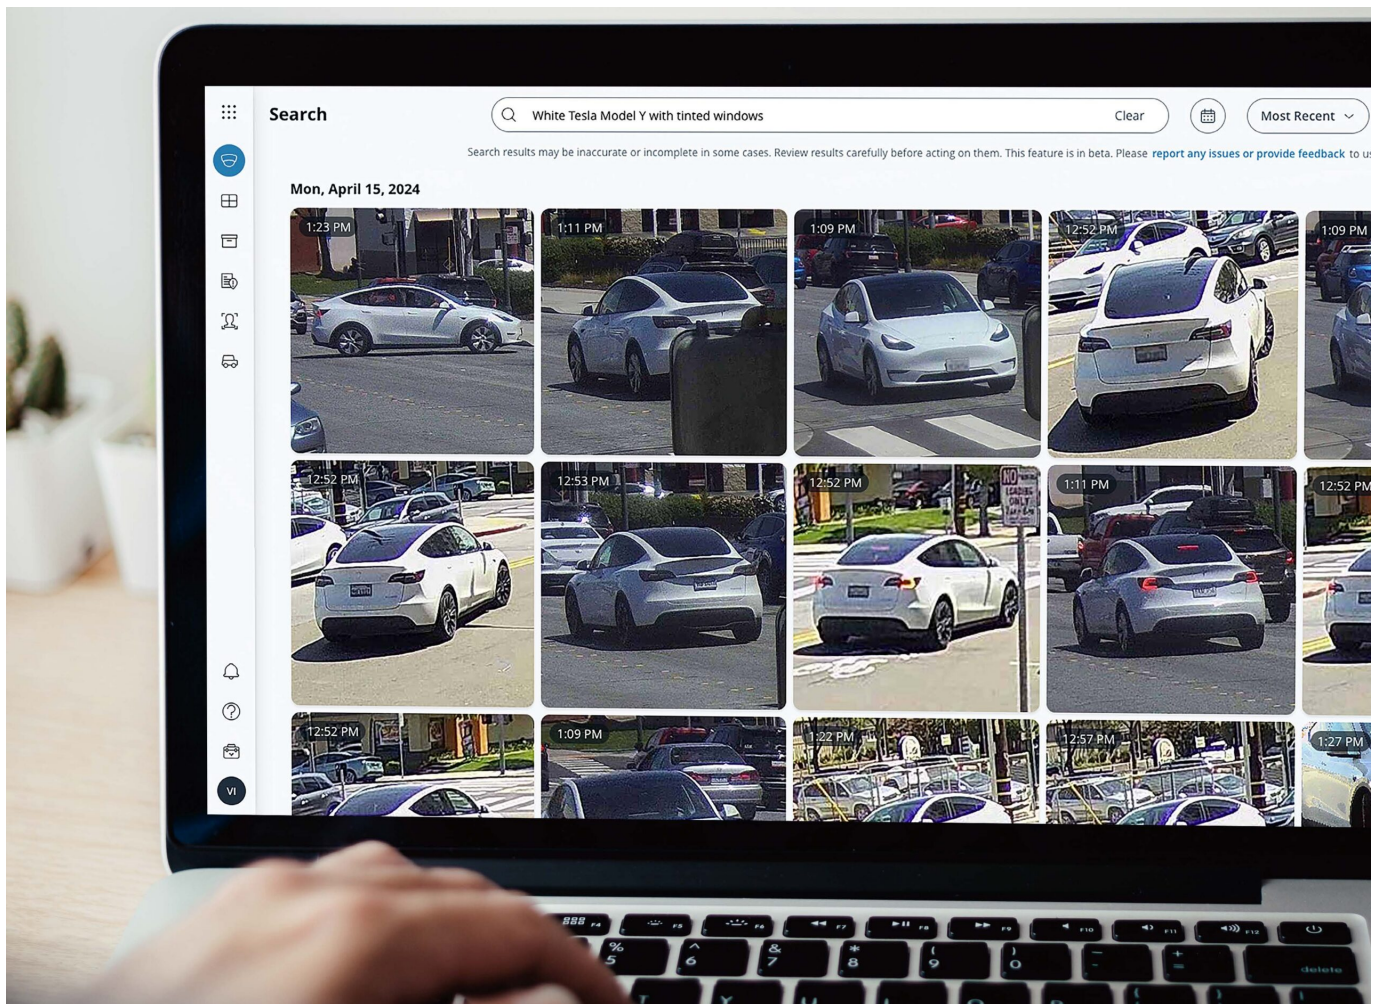
[Home](#) » [Verkada](#) » Verkada 0524 AI Powered Search Camera User Guide 

Contents

- [1 Verkada 0524 AI Powered Search Camera](#)
- [2 AI-Powered Search User Guide](#)
- [3 Overview](#)
- [4 How to Get Started and Create Actionable](#)
- [5 Optimal Query Structure and Suggested](#)
- [6 Below we include some example queries to help you get starte](#)
- [7 Limitations](#)
- [8 Query Moderation & Responsible Use](#)
- [9 FAQ](#)
- [10 Documents / Resources](#)
 - [10.1 References](#)
- [11 Related Posts](#)



Verkada 0524 AI Powered Search Camera



Specifications

- Product Name: AI-Powered Search
- Manufacturer: Verkada
- Website: www.verkada.com
- Email: sales@verkada.com

Setting up AI-Powered Search for Your Organization

1. Organization administrators need to enable AI-powered search.
2. Log in to Command and look for the prompt to Try the new AI-powered search.
3. Accept the AI-powered search Terms and Conditions to proceed.

Using AI-Powered Search

1. New queries search across all cameras by default.
2. You can organize searches by camera sites, camera type, popular footage under Live Detection, or recent searches.
3. Archive or flag search results for further action.
4. View, download, archive, and share relevant footage from the footage stream.
5. Note that the feature exists on the home camera screen page.

6. Clicking into a specific camera displays classic attribute search.

Optimal Query Structure & Best Practices for AI-Powered Search

AI-powered search can identify text, logos, and brands. It can assist in recognizing specific objects like people, vehicles, text, logos, and brands. However, it has limitations on identifying text and numbers. Follow best practices:

- Text recognition works best when text, logos, or brands are clearly visible.
- Understand the limitations when attempting text-matching searches.

FAQ

Q: What are the limitations of text recognition in AI-powered search?

A: The model does not use Optical Character Recognition (OCR) and may have limitations in identifying text and numbers. Text recognition works best when the text, logo, or brand is clearly visible.

AI-Powered Search User Guide

Overview

Verkada's AI-powered search leverages state-of-the-art large language and vision models that let users search for relevant camera footage of people and vehicles using everyday language, making investigations more intuitive and efficient. Users in Command are no longer restricted to a predefined list of search attributes of people or vehicles and can instead enter freeform text searches ("queries") directly into the search bar to get more granular results. Using one's own language, users in Command can now create searches like "person holding a mop" or "white Tesla model Y with tinted windows" to conduct highly detailed investigations.

How to Get Started and Create Actionable


Investigations with AI-Powered Search

Setting up AI-powered search for your organization

1. Organization administrators must first enable AI-powered search: i. have People and Vehicle Analytics enabled; ii. enable AI-powered search in the "Feature Manager" tab under "Privacy & Security" settings, pictured below under Step 1B and iii. offer the requisite permissions for "AI-powered search" to the desired group (i.e., "Site Viewer" or "Site Admin") within the "Roles & Permissions Customization" tab under "Org Settings," as pictured below under Step 1C.

Note: searches will not return results for Vehicle or People History if they are disabled. You can enable them by toggling them on in the "Feature Manager," under "Privacy."

Step 1B. for organization administrators enabling AI-powered search.



Step 1C. for organization administrators enabling AI-powered search.

Manage Analytics & Search Permissions

Select role needed for each permission type

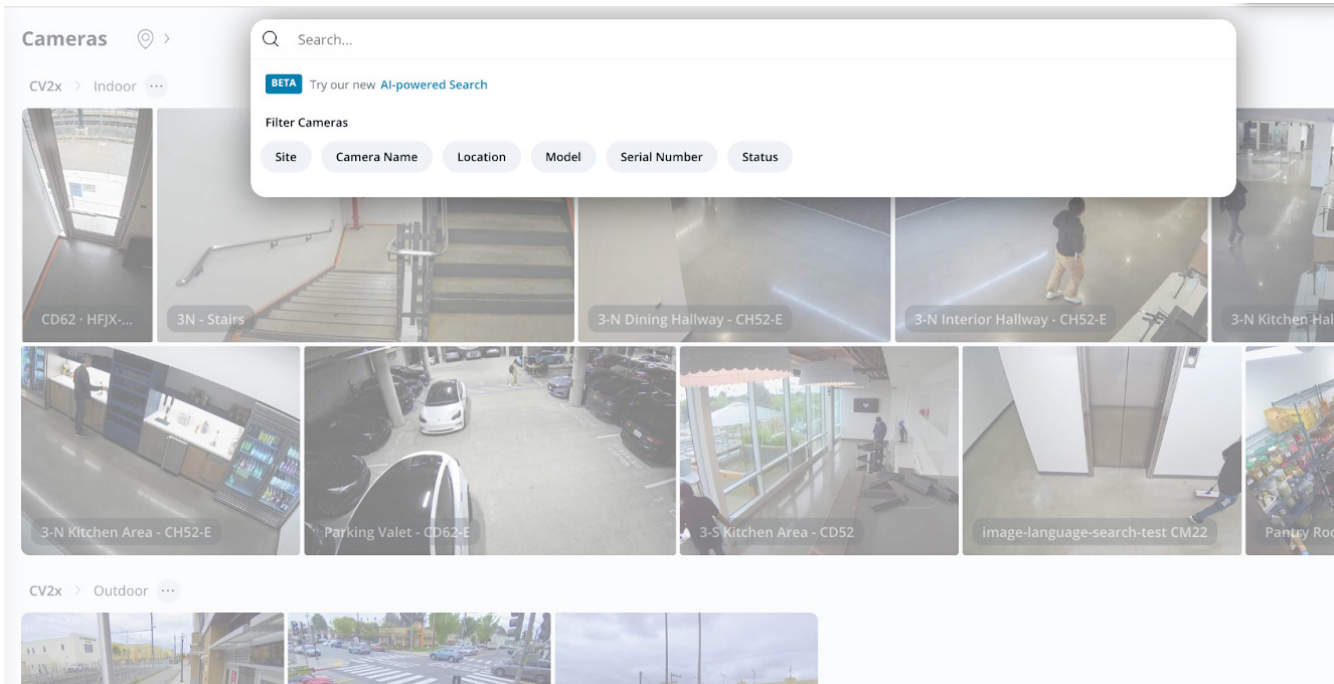
Cancel
Save

Permission	Site Viewer	Site Admin	Org Admin
View People Analytics	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
View People of Interest	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Manage People of Interest	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Use AI-Powered Search	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

[Restore Default Settings](#)

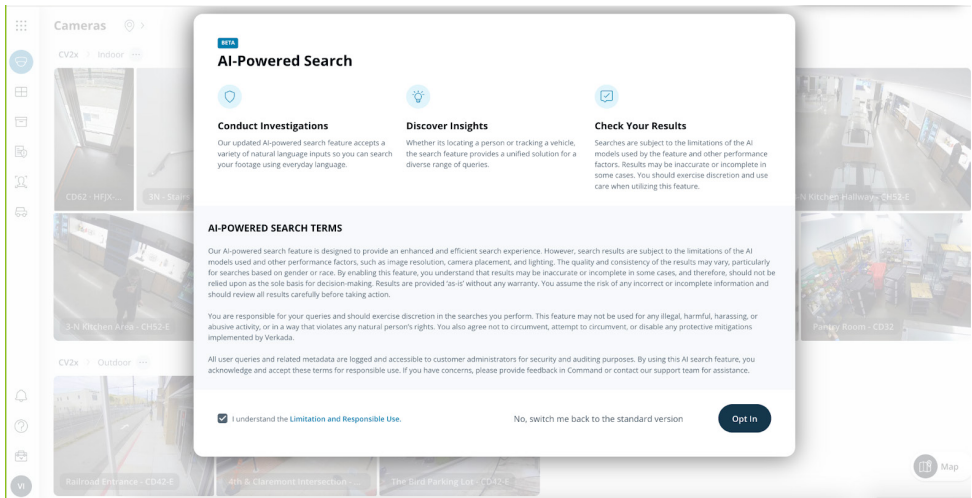
Using AI-powered search

2. Log in to Command and you'll see a prompt to "Try our new AI-powered search."



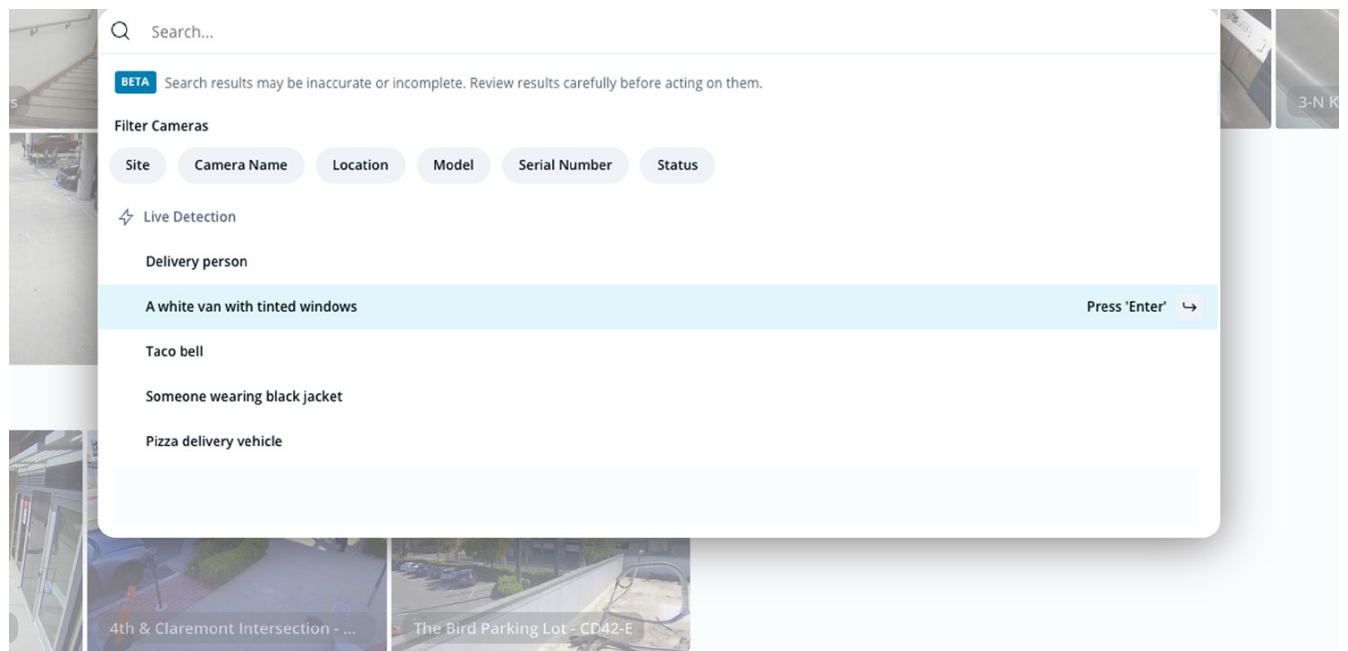
3. Once permissions are enabled and the user clicks "Try our new AI-powered search," the user will then be prompted to accept the AI-powered search Terms and Conditions (Terms). Please read through them carefully so you understand the capabilities and limitations of this feature.

Note: the user will only need to opt in once.

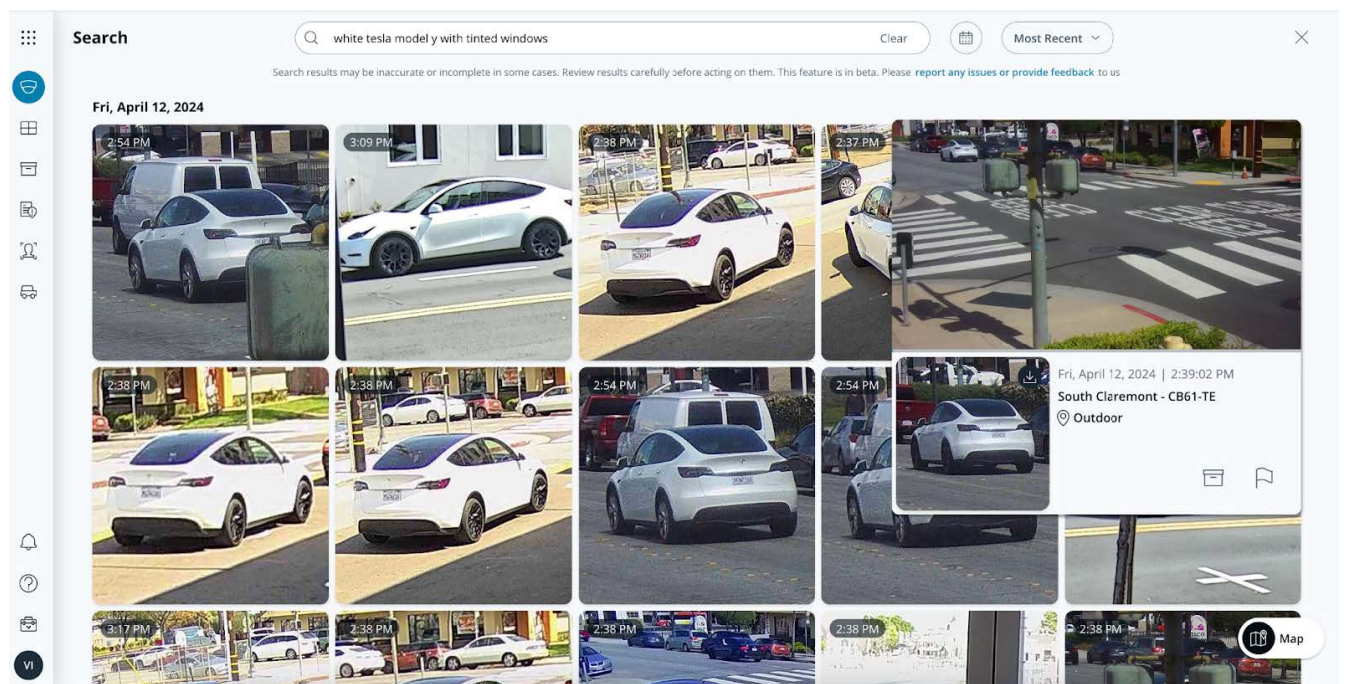


4. Upon accepting the Terms, you'll see a search bar that accepts any text—your blank canvas for writing queries in your own language, directly relevant to your investigations.
- New queries search across all cameras by default, but you can also organize your searches by camera sites, camera type, leverage searches of popular footage in your organization (under the "Live Detection" banner) or

revisit recent searches.

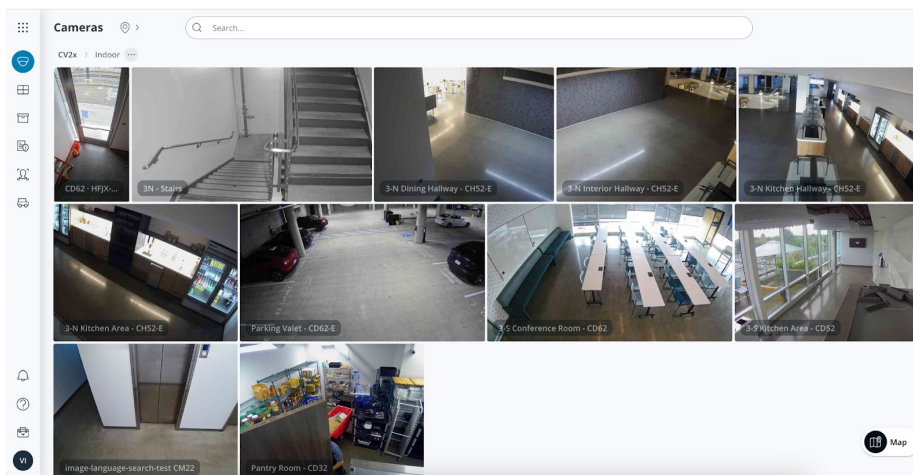


5. Hovering over a result gives you the option to archive the clip (via the small box icon) or flag the result and submit feedback.

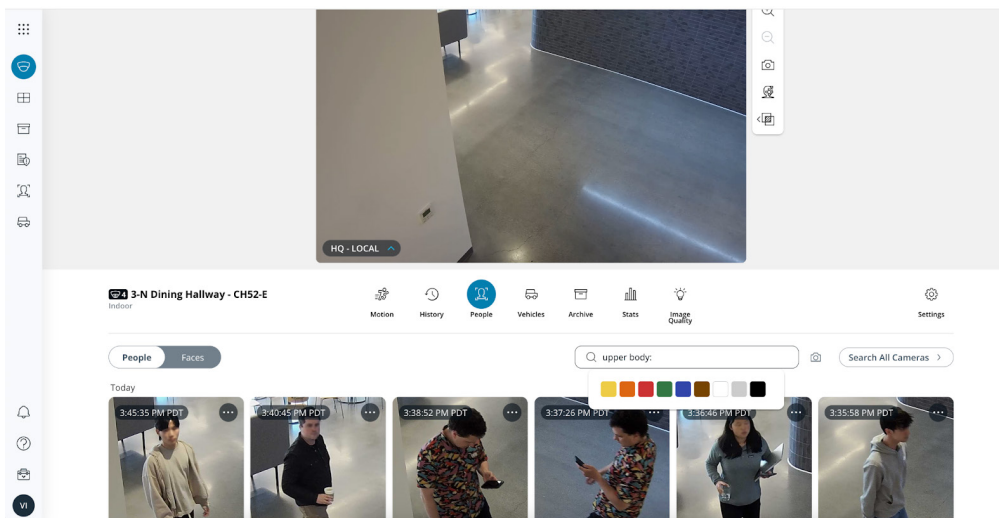


Clicking into the corresponding result brings you directly to the footage stream where you can download, archive and share the relevant footage.

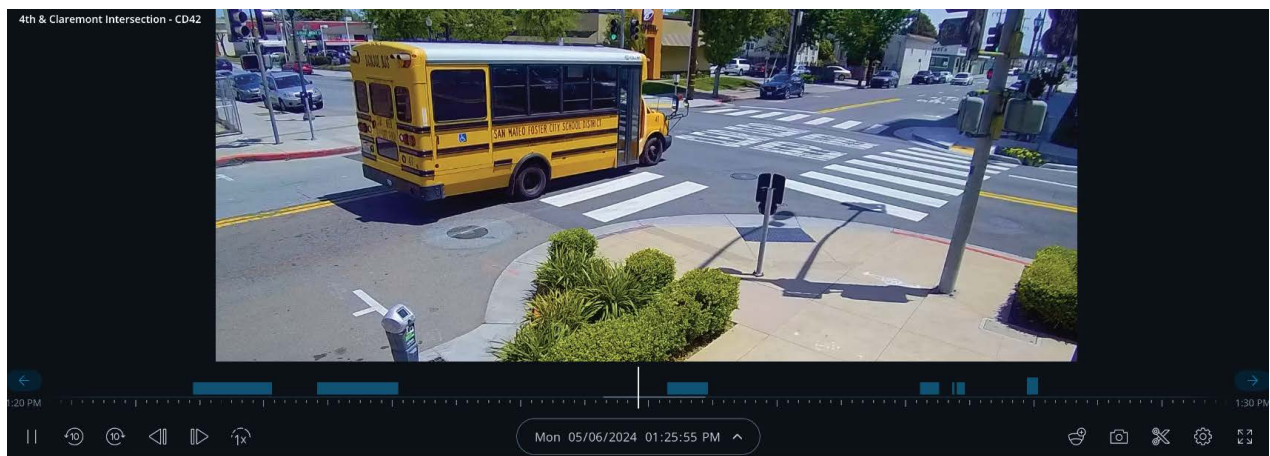
6. Note that this feature only currently exists on the home camera screen page. Clicking into a specific camera displays the classic attribute search, as depicted in the second picture below.



The initial search home page with AI-powered search in the first image. Upon clicking into a specific camera on the home page, the user will see our classic attribute search with a predefined list of criteria (like selecting shirt color)—not AI-powered freeform search, as depicted below.



7. Clicking into a search result brings the user to the specific footage in the history player.



Clicking into a result for the query “San Mateo School Bus” brings the user to the history player at the time of the requested footage, as depicted above. www.verkada.com

Optimal Query Structure and Suggested

Queries for Useful Investigations

You can read more about the underlying models that support this feature in our white paper here to understand why we suggest structuring queries in the method we highlight below. We recommend the following “best practices” when using AI-powered search:

Best practices

1. Use AI-powered search for highly-detailed attribute searching for people or vehicles, not standalone objects

AI-powered search leverages “hyperzooms,” (HZs), which are highly-detailed segments of a video frame that contain images of a person or a vehicle only, typically the most important parts of an investigation. If your query does not include a person or vehicle, results will be directed to HZs of people and vehicles that include the object (e.g., if searching for

“ski gear,” the search will return results of ski gear on a person or attached to a vehicle).

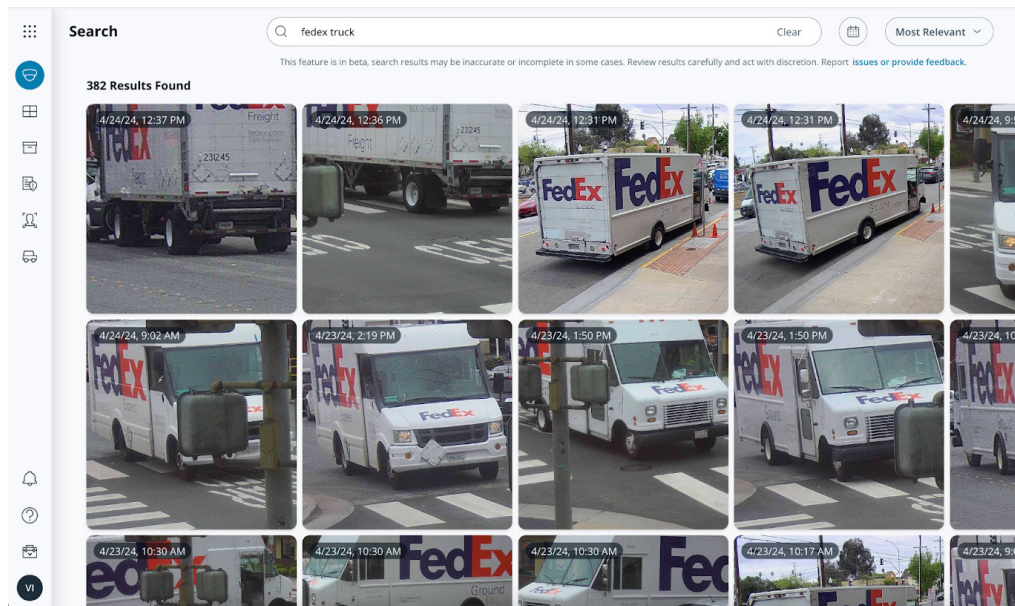
2. Use AI-powered search to filter further by time and space

AI-powered search includes a concept known as “parsing,” or the ability to include date, time and location parameters in a query. To further narrow a query using parsing, a user can add to the search time specifications like a date range

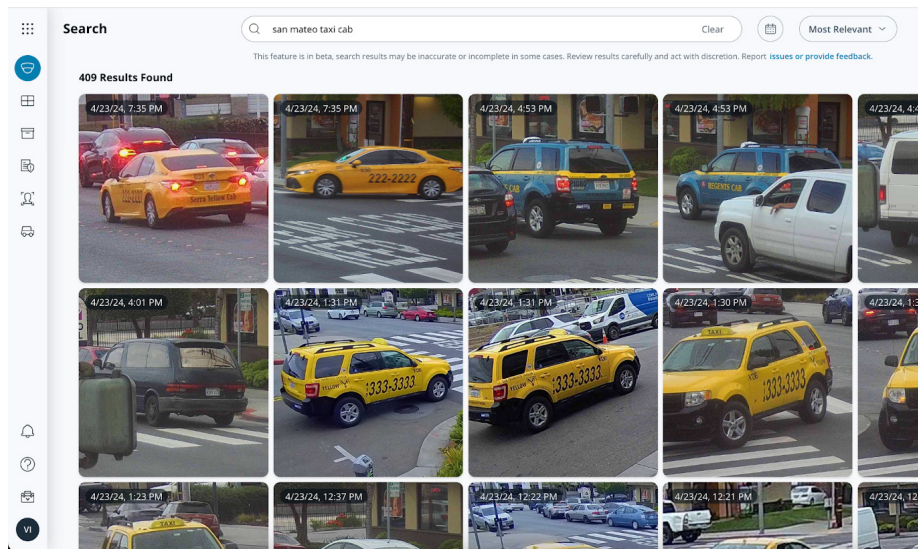
or specific day (e.g., “UPS truck between April 5th and April 6th,” or “UPS truck yesterday”), a time range (e.g., “person holding tote bag between 9am and 10am”), or a location tied to a camera (e.g., “people crowding inside the main lobby,” or “white van parked in the exterior parking lot”).

3. AI-powered search can be used to identify text, logos and brands

AI-powered search can assist you in identifying specific objects people hold or vehicles driven via recognition of sufficiently visible text, logos and brands, as pictured below. Note: this model was not trained using Optical Character Recognition (OCR) and therefore has limitations on identifying text and numbers. Please refer to the FAQ below for additional details.



Text recognition works best when the text, logo or brand is clearly visible as illustrated in the results above for “FedEx truck.”



A search for “San Mateo taxi cab” returns mixed results. The model has detected the words “taxi” and “cab,” but “San Mateo” isn’t clearly visible in any of the results. Customers should note these limitations when attempting text-matching searches.

With these best practices in mind, we suggest structuring queries as follows:

[person or vehicle] + [the specific action and/or descriptor of that person or vehicle] + [OPTIONAL “between” / “on” / “in” / “at” / “with”] + [OPTIONAL second or third descriptor] + AND/OR [OPTIONAL date / time / location]

Two example queries using the above structure:

“Person wearing a green scarf with red hat between 8am and 9am outside the main entrance”

“White Ford pickup truck with red bumper sticker yesterday”

Important notes on query inputs

- We’re actively working to better support the use of terms “AND,” “OR,” and “NOT,” but they currently do not produce the most accurate results. In lieu of these terms, we suggest using any of the other optional terms above.
- As the foundation model was trained on still images, it will perform less accurately on specific actions. Verbs such as “running” or “stealing” might produce less accurate results than “wearing” or “holding,” as “wearing” and “holding” can be more easily identified from still images than “running” or “stealing.”

Below we include some example queries to help you get started

Retail (anti-theft)

- “Person with duffle bag between 9 and 10 am”
- “Person holding [insert retail item]”
- “Person wearing Patagonia vest with scarf”
- “Person holding Gucci handbag”
- “Person wearing Louis Vuitton sweater”
- “Person carrying large shopping bag ”

General security

- “Person climbing over fence in [insert specific camera or site] on April 5th”
- “Person scaling wall”
- “Person looking at camera”
- “Person wearing balaclava”
- “Person with badge on chest”
- “Person walking near Sidewalk Closed sign”
- “Person near Amazon box”
- “People fighting each other”

General people search enhancements

- “Person wearing green hat with scarf”
- “Man with blue suitcase”
- “Person in spiderman costume”
- Man wearing New York shirt”
- “Man holding an apple”
- Person wearing a bike helmet”
- “Kid riding a bike”
- “Person with brown hair looking at their phone”
- “Person playing with their phone”
- “Person making a phone call yesterday”
- “Person walking a dog”
- “Person wearing a 49ers jersey”
- “Delivery person”
- “Person wearing Canada Goose jacket”
- “Person holding a Redbull can with sweatpants on April 7th”
- “Kids playing”
- “Person holding cat”

Education (student safety, rules compliance)

- “Person using a vape”
- “People in a crowd in the parking lot”
- “People on football field between 12 am and 4 am”
- “Cars in parking lot between 8pm and 12am on April 6th”
- “[Insert school district name] school bus”

Manufacturing (workplace safety, compliance)

- “Person wearing a hardhat with safety vest”
- “Person wearing a safety vest in [insert specific camera or site] on April 3rd”
- “Person carrying a box around 11 am”
- “Person driving a forklift”
- “Person near wet floor”


- “Person carrying crates”
- “Person stocking supply shelf in kitchen”
- “Person using [insert tool, machine or instrument]”

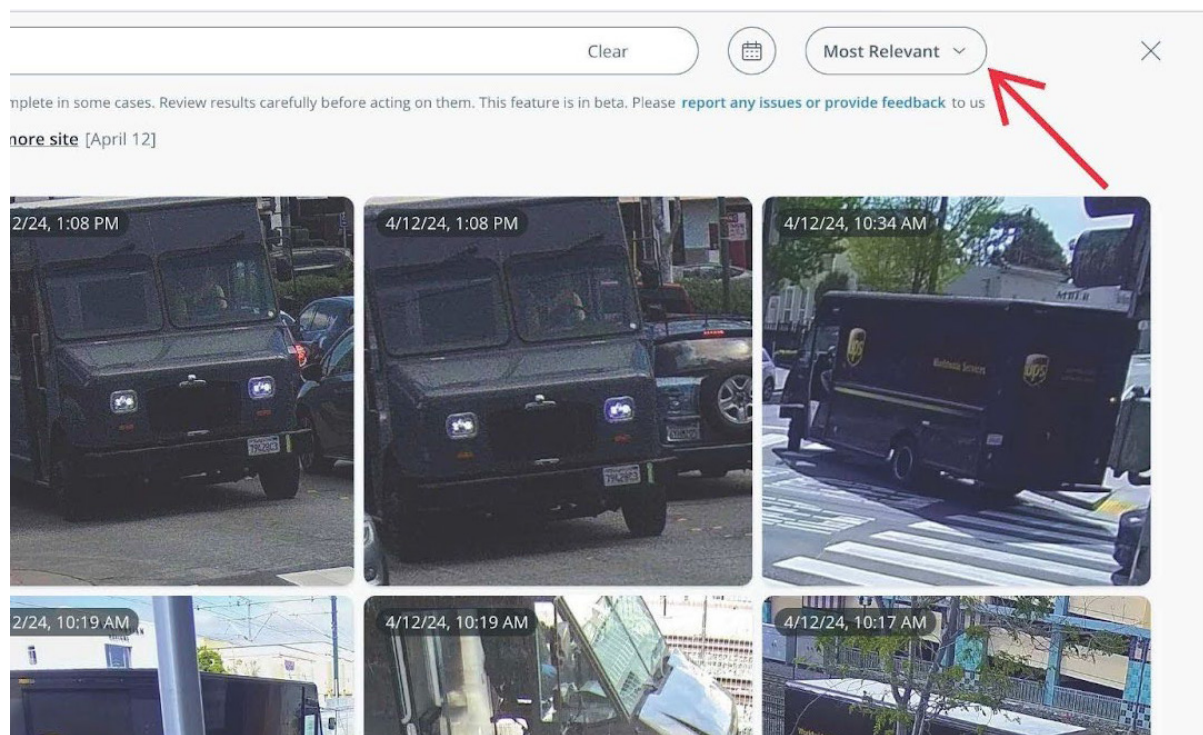
General vehicle search enhancements

- “White Ford Explorer with tinted windows”
- “FedEx truck at [insert specific camera or site]”
- “Gray BMW X7”
- “Black Tesla Model Y on April 4th”
- “Truck with roof rack”
- “Blue Porsche Taycan”
- “Most expensive car”
- “McLaren”
- “Porsche 911”
- “Gray 1960s sports car”
- “1995 Jeep”

Limitations

In this first iteration of AI-powered search, customers may experience incorrect or incomplete results, which we are working to improve, both for recall and precision.² We’ve opted for a model that has higher recall than precision.³ As a result, users will see all likely results of their investigation (inclusive of true positives and some false positives), and can determine the accuracy of their search results themselves. If we opted for a model with higher precision, some results identified by the system as false negatives would not be displayed to the user, which could be more costly than a false positive for an investigation. As designed, we put the customer in control to filter out false positive results for themselves. You can also boost precision by selecting the “Most Relevant” button instead of the “Most Recent” button, as depicted below:





Two different sets of results from the same query “UPS truck yesterday.” In the first set of results with “Most Recent” toggled on, the model displays the most recent results from all cameras that could reasonably resemble a UPS truck (i.e., higher recall than precision). In the second set of results when the user selects “Most Relevant,” the model discounts more recent results for results with higher precision and produces fewer false positives.

- In machine learning modeling, “precision” is the percentage of true positive results out of the entire batch of “positive” (inclusive of true and false positives) returned results. It’s expressed formulaically as $[\text{Precision} = \text{True Positive Results} / (\text{True Positive Results} + \text{False Positive Results})]$. It’s a measure of the model’s accuracy in identifying positive results. “Recall,” on the other hand, measures the percentage of true positives that are correctly identified. It is expressed formulaically as $[\text{Recall} = \text{True Positive Results} / (\text{True Positive Results} + \text{False Negative Results})]$. Whereas precision measures the overall rate of correct predictions, recall is helpful to identify the rate of false negatives. See here for a depiction of precision and recall.
- Note that with Large Language Models (LLMs), recall and precision statistics will differ depending on the measured sample sizes. Our model’s average recall is greater than its average precision.

Query Moderation & Responsible Use

Unlike our filter-based attribute search, which has a predefined set of descriptors for narrowing search results, our AI-powered search lets customers search their footage for a wide range of attributes using their own words. Because our foundation model has been trained on publicly available data from the Internet, results may still be inaccurate, inappropriate or offensive. We have built query moderation into our platform to help reduce the risk that our AI-powered search is used maliciously or in ways that may be harmful. Moderation is a critical safeguard for this powerful feature.

At the same time, we implemented moderation in a way that isn’t overly restrictive to the point that it compromises the feature’s usability. We also leverage industry-recognized practices, including open source data and OpenAI’s moderation APIs, to make moderation more effective. Striking the right balance between respectful and useful searches is paramount for us.

Our approach to moderation

Just as we’ve chosen to use a publicly available model for our AI-powered search feature, we developed our query moderation using publicly available practices for effective moderation. OpenAI, for instance, has published guides and blogs for developing suitable moderation techniques.⁴ We further augmented these capabilities with

additional proprietary protections, which include our own list of prohibited search categories:

- Information about the race, ethnicity or nationality of a person
- Information about the educational qualifications or religious beliefs of a person
- Subjective descriptions of people (e.g., attractive, ugly, wealthy)
- Inferred content from an image (e.g., “coolest person in the world”)
- Sexual content or innuendo
- Names of specific people (i.e., public figures)
- Inappropriate or offensive descriptions of people or objects (e.g., “dumb person”)
- Slurs equating people with animals (e.g., monkeys)

We also tapped into open source data to generate a list of banned text strings, in languages with latin characters like English, Spanish, French, German, etc.⁵ When words or phrases in a query appear on the list, we reject the search and mask any results—a process known as string-based matching. We also complemented this technique with OpenAI’s moderation API to cross-check queries that create harmful or biased results.

We also hold frequent engineering “bug bashes,” analyze thousands of our own engineering test queries (i.e., classify thousands of queries as “appropriate” or “inappropriate”) to modify our moderation rules and optimize them for usability and to avoid bias. We can update our “blocklist” in minutes and thereby continuously improve our moderation techniques.

- See blog: <https://openai.com/blog/using-gpt-4-for-content-moderation>. For more detailed, technical approaches to moderation that we leveraged, see: https://cdn.openai.com/papers/DALL_E_3_System_Card.pdf and https://cdn.openai.com/papers/GPTV_System_Card.pdf.
- Note: currently our AI-powered search only supports queries in languages with Latin characters.

FAQ

1. Why did we develop AI-powered search?

Beyond common characteristics such as the presence of a backpack or the color of a car, many of our customers need the ability to search for people and vehicles across a much wider range of descriptive features. This capability is important for security and physical operations professionals, as it enables them to search through their footage quickly with a much more expansive set of descriptors than before. It also addresses operational issues across a variety of industries—from retailers looking to identify shoplifters or manufacturing customers looking to support workplace safety (e.g., “person driving a forklift”).

2. How did we develop the AI-powered search capabilities?

When developing this feature, we employed a two-pronged approach. First, we leveraged a publicly available model with large language and large vision functionalities. This model enabled direct comparisons between text and images by training a multi-billion parameter neural network to bring related images and texts closer together while pushing unrelated ones apart. The foundation model also provided the basis for image classification and retrieval, allowing users to search for images using natural language.

Our approach further builds upon the foundation model in several important ways. We have, for instance, built a scalable, logically separated vector storage and retrieval system that processes customer video data before the search query is performed (in advance in a cache) to enable faster queries in real time, (e.g., no need to run the customer’s video footage through the foundation model each time a query is run). This allows our AI-powered search to index and retrieve relevant footage for our customers quickly and at scale. See [here](#) for more

information about our foundation model and how we've improved upon it.

3. What data do we send to OpenAI and what do we get back in return?

We send user search queries to OpenAI and receive back in return a "pass" or "fail" response. When using it for query parsing, we send the query to OpenAI and get it back parsed with time and camera/site information extracted from the original query which helps provides more accurate results.

4. Does Verkada use my queries to develop or improve its products or services?

We use queries only to provide the AI-powered search service, and will not use queries to develop or improve our model or other offerings without the customer's consent. If a customer submits feedback on specific results from AI-powered search, however, we may use that feedback to improve our products or services or to enhance the customer experience.

5. Can I use AI-powered search for all Verkada products?

No. As described above, there are feature limitations tied to people and vehicles (standalone non-person/vehicles are not recognized). We also deploy query moderation to reduce the risk that our AI-powered search is used in malicious or in harmful ways.

6. Can I use AI-powered search to find anything?

No. As described above, there are limitations as this search is tied to people and vehicles (and not standalone non-person, non-vehicle objects).

Also, we're employing query moderation to reduce the risk that our AI-powered search is used maliciously or to promote harm.

7. Can AI-powered search be used to identify text, numbers, logos or brands on specific objects?

In order for AI-powered search to identify specific text, numbers, logos or brands on objects, two criteria must be met: i. the object needs to be tied to a person (e.g., a person is holding a bag with a clearly visible brand) or a vehicle (e.g., text on the side of a truck or bus) and ii. the text, number, logo or brand that you're trying to identify must be sufficiently visible for the camera to detect. Ideally, the text or logo should be large, clearly visible and the camera should have a high resolution. If these criteria are not met, then the results will be less accurate. Note: this feature was not trained using Optical Character Recognition (OCR), so it cannot be used reliably to read license plates, barcodes, SKU names or other alphanumeric codes.

8. Is there an additional charge to access AI-powered search?

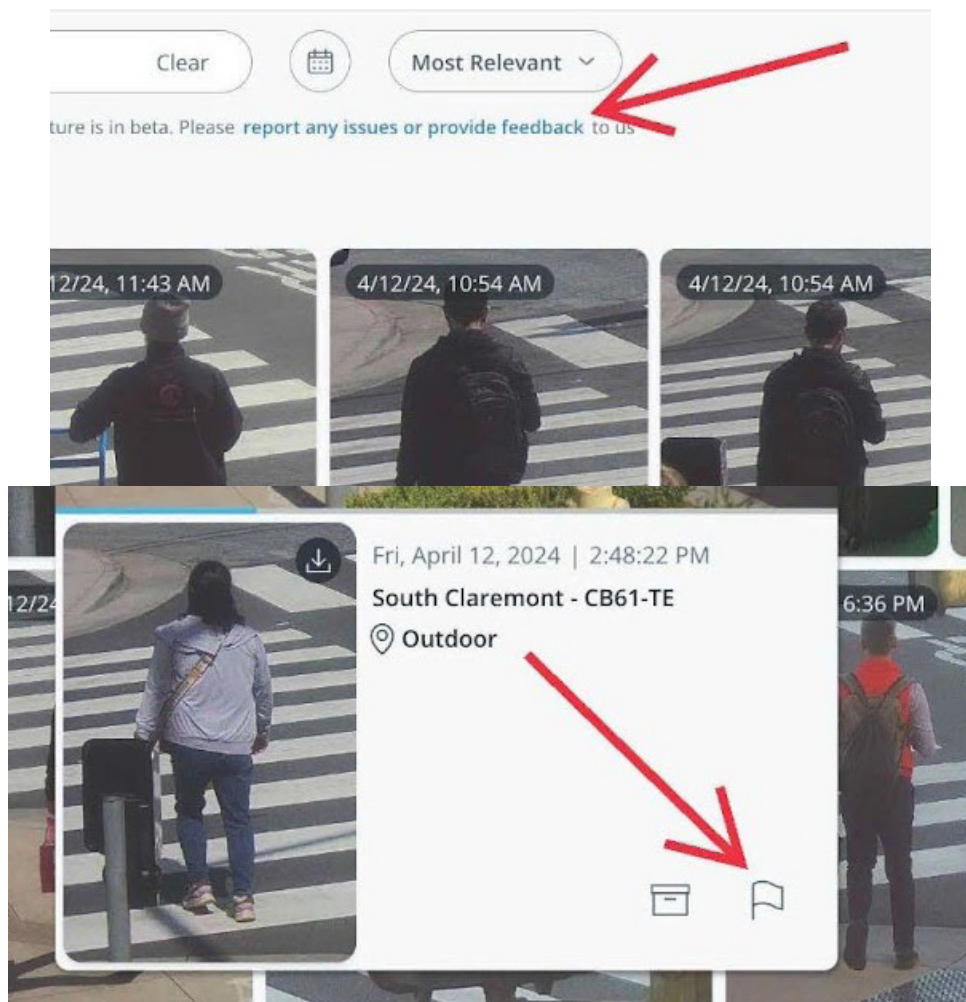
No. All camera customers with Command licenses receive access to AI-powered search without additional charge.

9. Does AI-powered search work on the Command Mobile app?

No. In its current state, AI-powered search works on the Command web browser. Customers can access AI-powered search from their computers or mobile phone using a web browser.

10. What if I see inappropriate or inaccurate results?

Please report your issue(s) to us using either option depicted below. We appreciate your feedback in helping us to continually improve our model and offering:




You can submit feedback on overall results and experience using AI-powered search via the button in the left-side image, or submit feedback on specific results using the flag button by hovering over a specific result.

www.verkada.com

sales@verkada.com

Documents / Resources



[Verkada 0524 AI Powered Search Camera](#) [pdf] User Guide
0524, 0524 AI Powered Search Camera, AI Powered Search Camera, Search Camera, Camera

References

- [User Manual](#)

[Manuals+](#), [Privacy Policy](#)

This website is an independent publication and is neither affiliated with nor endorsed by any of the trademark owners. The "Bluetooth®" word mark and logos are registered trademarks owned by Bluetooth SIG, Inc. The "Wi-Fi®" word mark and logos are registered trademarks owned by the Wi-Fi Alliance. Any use of these marks on this website does not imply any affiliation with or endorsement.