

TRADING KNOWLEDGE LIBRARY

What the Machine Can Actually See

Market Analysis



Trading Papers

What the Machine Can Actually See

An assistant reading your chart, and the four things it is not doing

Published by	Trading Papers
Series	The trading knowledge library
Edition	First edition · 2026-08-27
Format	A4 · digital

Educational content only. Nothing in this document is financial advice, no part of it is a recommendation to buy or sell anything, and it knows nothing about your circumstances. Trading carries risk, including the loss of the amount invested.

© 2026 Trading Papers. All rights reserved.



Trading Papers

Contents

01	What is actually happening	4
02	The four things it does well	5
03	The four things it cannot do	6
04	The fourth one, which costs money	7
05	What a live connection actually adds	8
06	What it needs to be told	9
07	Asking something that has an answer	10
08	What a good answer looks like	12
09	Checking it in two minutes	12
10	Keeping the record	14
11	Where it genuinely saves time	15
12	Where it costs you money	16
13	The one it cannot check for you	17
14	What it changes and what it does not	18

01

What is actually happening

An assistant looking at your chart is doing exactly one thing. It is turning a picture into a description.

What you think you gave it	What it received
Your chart	an image, or a table of numbers
The market	the part of it inside the frame
Your question	words, with no account behind them
A decision to make	a request to produce text

What actually crosses the gap when an assistant looks at a chart. The left column is what you believe you handed it; the right is what arrived. Every capability and every failure in this paper is a consequence of that second column.

That is not a diminishment and it is not a warning. It is the mechanism, and every useful thing in this collection and every expensive mistake in it are consequences of that one line.

Read the right column of that table again, because the gap between the two columns is where the trouble lives. You believe you handed it your chart; it received an image. You believe you asked it a question about the market; it received words, with no account and no history behind them.

The first paper in AI Trading. It contains almost no instructions for any particular tool, on purpose: menus move and this does not.

02

The four things it does well

Start with what works, because it is genuinely useful and the rest of this paper is about limits.



The four things it does well, ranked by how much time each actually saves against doing it yourself. Naming and counting are where the hours go; drafting is the one people demonstrate and the one that saves least, because a draft still has to be checked line by line.

It names things. Shown a chart, it will tell you what is on it in ordinary words — a range, a gap, a sequence of higher lows — and it does that in seconds, correctly, most of the time.

It counts, and this is the one to keep. A human counting swings across two hundred bars loses their place around bar sixty and starts guessing; the machine does not get bored at bar sixty. Counting is also the single most checkable thing it produces, which makes it the safest thing to ask for.

It compares, which is the same skill applied twice. And it drafts: given a description of what you do, it will write it out as rules. That last one impresses people the most and saves the least, because a draft you have not checked line by line is not a method, it is a suggestion.

03


The four things it cannot do

Now the other side, and the reason this paper comes first in the collection.

It cannot	Because	What that looks like
See off the screen	the frame is the whole world	a level from last year it never mentions
Know your account	nobody told it the balance	position sizes that would ruin you
Know what happens next	nothing can	a forecast delivered in the same tone as a fact
Decline to answer	it is built to respond	a confident reading of noise

The four blind spots. The first three are limits of information and are obvious once stated; the fourth is a limit of behaviour, is not obvious at all, and is the one that empties accounts.

The first three are limits of information and are almost obvious once written down. It cannot see past the edge of the frame, because the frame is the entire world it was handed. It does not know your balance, because nobody told it. It cannot know what happens next, because nothing can.



The first blind spot, drawn. The assistant was shown the marked window and described a clean advance off support, which is a correct description of the window. The level that actually governs this instrument is above the frame, and it was never in the conversation.

That figure is the first blind spot doing its damage quietly. The assistant was shown the marked window and described a clean advance off support, which is an accurate description of what it was given. The level that actually governs this instrument sits above the frame, and it never came up — not because the model missed it, but because it was never there.

04


The fourth one, which costs money

The fourth is not a limit of information. It is a limit of behaviour, and it is not obvious at all: the assistant cannot decline to answer.

	There is something to see	There is nothing to see
It answers	A useful reading. This is the case people remember.	A confident reading of noise. Reads identically.
It says so	Never happens. It is built to produce a reply.	Never happens either, which is the whole problem.

Why the fourth blind spot is the expensive one. The two cells on the right are indistinguishable in the reply you receive: both are fluent, both are specific, and only one of them is about anything. Nothing in the wording tells you which you are holding.

Look at the right-hand column of that grid. When there is genuinely nothing on the chart, you do not receive a reply saying so. You receive a reading — fluent, specific, and about nothing. It is indistinguishable in tone from the reading you get when there is something to see.



How a wrong answer gets its authority, in the order the reader experiences it. Nothing in this sequence involves the model being unusually wrong; it is the ordinary case, dressed in prose that has no way of sounding uncertain.

Nothing in that sequence requires the model to be unusually wrong. It is the ordinary case. Prose has no way of sounding uncertain unless it is written to, and a clear paragraph reads as a considered one whether or not anything considered it.

That is the whole of the difference between this tool and a colleague. A colleague looks at a quiet chart and says there is nothing here. That sentence is the most valuable one in analysis, and it is the one sentence the machine will not produce unprompted.

05

What a live connection actually adds

You can wire an assistant to a charting platform so that it reads live prices rather than a screenshot you pasted. It is worth doing, and it is worth being exact about what it changes.

Connecting it gives you	Connecting it does not give you
Current prices instead of a screenshot	any knowledge of your account
The bars, as numbers	judgement about which bars matter
Indicator values it can read exactly	a reason to trust the indicator
Several timeframes on request	the ability to know which one to use

What a live connection to a charting platform changes, and what it leaves exactly where it was. The right column is the honest half: connecting the tools moves the data, and moves none of the limits.

The left column is real. Live numbers beat a screenshot: no misread axis, no stale prices, exact indicator values rather than an estimate off a picture, and several timeframes on request instead of four screenshots.

The right column is the half that gets skipped in every tutorial. Connecting the tools moves the data and moves none of the limits. It still cannot see beyond the window it pulled, still does not know your balance, still cannot tell which timeframe is the right one, and still cannot decline.

This paper deliberately does not tell you which menu to click. Any such instruction would be wrong within a year, and the four limits above will not have moved.

06


What it needs to be told

A connection hands it the numbers. It still knows nothing about you, and the sentence you type is the only place that can come from.

It does not know	Unless you say	Or the answer is built on
Which instrument this is	the name	a generic chart
Which timeframe	the bar width	a guess from the shape
What you are deciding	the actual question	whatever it assumes
What you already believe	your reading	no context for disagreeing
How long you can hold	your horizon	advice for someone else

Everything the assistant does not know about your situation unless the sentence you typed contains it. None of this is inferable from the chart, and every row of it changes what a sensible answer would be.

Read the right column. Every row is a real answer the model will produce in the absence of the information — not a refusal, not a question back, but a reply built on a default. Ask about an unnamed chart and you get a reading of a generic chart, delivered as though it were about yours.



What a first message should carry, in the order that matters. The last step is the one people leave out, and it is the one that turns the reply from a description into something addressed to a decision you are actually facing.

The fourth step is the one people leave out and the one that changes the most. Saying what you already think gives the model something to disagree with, and disagreement is the only mode in which it is useful as a check rather than a mirror. Left to itself it will elaborate whatever direction your question leaned.

None of this is a magic phrasing, and no wording gets you a better model than the one you have. It is the ordinary business of telling somebody the things they would need to know, which feels unnecessary only because the reply arrives whether or not you did.

07

Asking something that has an answer

Between a session that saves you an hour and a session that wastes one, there is a single difference, and it is not the model.

Unanswerable	Checkable
Is this bullish?	How many higher lows in the last twenty bars?
Where is support?	Which price did it stop falling at more than once?
Should I buy here?	What would have to be true for this setup to fail?
Is this a good strategy?	How many trades did this rule produce last year?

The single largest difference between a useful session and a wasted one. Left is a question whose answer cannot be checked; right is the same curiosity, asked so that the reply is either right or wrong and you can tell which.

It is whether the question has a checkable answer. Ask whether a chart is bullish and you get a paragraph you can neither confirm nor argue with. Ask how many higher lows there are in the last twenty bars and you get a number, which is either right or wrong, and you can tell which by looking.

Where a question splits. The left branch produces a paragraph you cannot argue with and cannot verify; the right produces a number you can look at the chart and confirm in under a minute. Both replies will be equally confident.

Both branches of that fork produce equally confident replies. Only one of them produces something you can catch being wrong, and being able to catch it is the entire value of asking.

The fourth row of that table is worth reading twice. Asking what would have to be true for a setup to fail is the one question that turns the machine's fluency into an asset: it is good at generating the failure cases you did not think of, and generating them costs you nothing.

08

What a good answer looks like

Once the question is right, the reply can still be worth little or a lot. There is a shape to the good ones.

- Names** ● what is on the screen, in ordinary words
- Counts** ● how many, over how many bars
- States** ● the rule it used to count that way
- Admits** ● what it could not see, and what would change it

What a good answer contains, from the bottom up. A reply that stops at the second rung is a description; one that reaches the fourth is something you can act on and, more to the point, something you can catch being wrong.


A reply that names things is a description, and a description is where most sessions stop. A reply that counts is better, because it can be wrong in a way you can see. A reply that states the rule it counted by is better again, because now you can apply the rule yourself somewhere else and find out whether it holds.

The top rung is the one to ask for by name, since it rarely arrives unrequested: what could you not see, and what would change this reading. A model asked that question answers it well. A model not asked it will not volunteer that the window was too narrow.

09


Checking it in two minutes

Everything above is preparation for this section, which is the practical core of the paper.



The two-minute check, in the order that catches the most for the least effort. The first step alone finds the majority of bad answers, and almost nobody does it because a fluent reply does not feel like it needs checking.

The first step is almost the whole method. Take one number it gave you and count that one thing yourself. It takes thirty seconds and it is the step nobody performs, because a fluent reply does not feel like something that needs checking.



What each check catches, from a session of forty questions put to an assistant about live charts. The first bar is the finding worth keeping: most of what goes wrong is arithmetic, not judgement, and arithmetic is the cheapest thing in the world to verify.

That bar chart is the finding worth keeping from a session of forty questions put to an assistant about live charts. Most of what goes wrong is arithmetic rather than judgement — a miscount, a level read off by a bar, a comparison that mixed two timeframes — and arithmetic is the cheapest thing in the world to verify.

The last step is the discard rule, and it is worth stating as one. If nothing would make the reading wrong, the reading is not about the market. Delete it and ask something else.

10

Keeping the record

Checking one answer tells you about that answer. Keeping the answers tells you whether the tool is any use to you on this market, which is a different and larger question.

Keep	Checkable later?	What it later tells you
The exact question you asked	yes	whether the question was answerable
The counts it produced	yes	whether it counts reliably on this market
The rule it said it used	yes	whether the rule holds elsewhere
Its verdicts and forecasts	no	nothing; do not store these

What to keep from a session, and why each line earns its place. The middle column is the test: anything you cannot check later is not worth storing, and anything you can is worth storing whether or not it turned out well.

The middle column is the whole filter. Store the question, the counts and the stated rule, because in a month you can go back and find out whether the counts were right and whether the rule held. Do not store its verdicts: there is nothing to compare them against later, and a folder of confident forecasts is a folder of things you cannot grade.

Counts of visible things  held up on re-checking

Levels it named 

Comparisons across timeframes 

Judgements about what happens next not checkable at all

How often each kind of claim survived being checked, from the same forty-question session. The gap between the top two rows and the bottom one is the argument of this whole paper in one figure: it is reliable at the mechanical and unreliable at the interpretive, in the same voice.

That figure is what a record gets you, and it is the argument of this paper compressed into

four bars. Counts of visible things survived re-checking nine times in ten. Judgements about what happens next could not be checked at all — not because they were wrong, but because there is no procedure that would have told you either way.

Both of those arrived in the same voice, in the same session, from the same tool. The record is how you learn to weight them differently, and there is no way to learn it from any single reply.

11

Where it genuinely saves time

With the limits established, the gains are easier to state honestly.

Task	By hand	With an assistant
Counting swings across 200 bars	twenty minutes, error-prone	seconds, then you verify one
Reading four timeframes in order	fifteen minutes	two minutes
Turning a habit into written rules	an evening	a draft in ten minutes
Comparing this week to a month ago	you do not bother	you now bother

Where the time actually goes back into your day. Every row is a task that is mechanical, bounded and checkable — which is the same list, seen from the other side, as the tasks where a mistake shows up immediately.

Every row on that table is mechanical, bounded and checkable, which is not a coincidence. Those are the same properties, seen from the other side, that make a mistake show up immediately.

The last row is the one people underrate. There are comparisons you have never made because they were tedious rather than difficult — this week against a month ago, this instrument against four others, this rule on a different market. The tedium was doing real damage to your work, and it is the tedium that has gone.

12

Where it costs you money

And the other column, which is shorter and more expensive.

Ask it for	What comes back	The damage
A position size	a plausible number	it does not know your balance
A prediction	a direction, stated calmly	nothing can know this
Whether a strategy is good	an assessment	it has not tested anything
A reading of a quiet chart	a reading	there was nothing there

Where it costs instead. Each row is the same mechanism: something that requires knowledge the model does not have, answered anyway because it always answers.

Every row is the same mechanism. You asked for something that requires knowledge the model does not have, and it answered anyway, because answering is what it does.



Six bars with nothing in them: no trend, no level touched twice, no range worth naming. Asked what it sees here, an assistant produces a paragraph — that is what it is for — and the paragraph will be about the third bar, because the third bar is marginally the largest.

That figure is the last row, drawn. Six bars with nothing in them: no trend, no level touched twice, no range worth naming. Asked what it sees, an assistant will produce a paragraph, and the paragraph will be about the third bar — because the third bar is marginally the largest, and marginally the largest is the closest thing to a pattern available.

You will not be able to tell that paragraph from a good one by reading it. That is why the check exists, and why the check is a count rather than a second opinion.

13

The one it cannot check for you

Every check so far catches a wrong answer. None of them catches the failure that costs the most, because it happens when the answer is agreeable rather than wrong.

	You were right	You were wrong
It agreed with you	Confirmed, and you learned nothing you did not have.	Confirmed, and now two of you are wrong.
It disagreed	You checked, held your reading, and it is stronger.	The only cell where the tool paid for itself.

The trap that survives every other check in this paper. The row is whether the assistant agreed with you; the column is whether you were right. Three of these cells teach you something and one of them teaches you nothing, and it is the one that feels best.

Read the top row. When it agrees with a reading you already had, you have learned nothing either way — if you were right you knew already, and if you were wrong there are now two of you. Agreement feels like verification and is not: the model had no independent access to the market, only to your framing of it.

The bottom-right cell is the only one where the tool paid for itself, and it is the one people argue with. Which is why the useful habit is the one from the prompt section: say what you already think, and ask it to make the case against. A model asked to disagree does it well, and the disagreement costs nothing to read.

14

What it changes and what it does not

It is worth sorting the effects before deciding how much any of this matters to you.

It changes	It does not change
How long the looking takes	whether the method has an edge
How many charts you can cover	whether you can hold a position
How quickly a habit becomes a rule	whether the rule survives testing
Whether you bother to check the wider frame	what the costs come to

What an assistant changes about the work and what it leaves untouched. The right column is every part of trading that decides whether you make money, which is worth noticing before deciding how much this matters.

The left column is real and is mostly about speed and coverage. The right column is every part of trading that decides whether you make money: whether the method has an edge, whether you can sit through the bad stretch, whether the rule survives being tested, what the costs come to.

- Sound** ● read the screen back to you, faster than you can
- Sound** ● count what you would have counted by hand
- Sound** ● write down the rule you are already using
- Unsound** ● decide whether the rule is any good
- Unsound** ● size the position, or press the button

The honest place for it in a working method, one rung at a time. Every rung above the third asks the model for something it has no way of supplying, and the fall from the top rung is the whole subject of the third paper in this collection.

The ladder puts that in one picture. The three sound rungs are all versions of reading and

writing down; the two unsound ones both ask the model to supply judgement it has no material for. Every failure in this collection is somebody standing on rung four or five.

15

What this does not settle

Three limits, and then the short version.

The limit	What it means here
Models change	the blind spots are structural; the details are not
Tools change faster	any menu named here would be wrong by spring
Nothing here is a test	a claim about an edge still has to be counted

Three limits on this paper itself, in the spirit of the rest of the library. The middle one is the reason there is almost no software instruction in these pages.

Models change, and the four blind spots in this paper are structural rather than a description of any current one. A frame is a frame; an account balance is not in the picture; the future is not knowable; and a system built to produce a reply will produce a reply. None of that is a version number.

Tools change faster still, which is why there is almost no software instruction in these pages. The second paper in this collection covers what to do with the assistant once it is looking at a chart, and the third covers rules, testing and the point at which people hand over more than they should.

Term	As used here
Assistant	a language model you are showing a chart to.
The frame	the window of the chart it was given, and nothing beyond it.
Checkable	a claim you can confirm or refute by looking.
Fluency	how finished the reply sounds; unrelated to whether it is right.
Verification	counting one of its numbers yourself before using any of them.

The terms this paper uses in a fixed sense. They carry the same meaning through the two papers that follow.

The useful summary is three sentences. An assistant reading your chart is converting a picture into a description, so it can name, count, compare and draft — and it cannot see past the frame, cannot know your account, cannot know the future, and cannot decline to answer. The fourth of those is the expensive one, because a confident reading of nothing is written in exactly the same prose as a good reading of something. So ask questions with checkable answers, and count one of its numbers yourself before you use any of them.

Educational content only. Nothing in this document is financial advice, and no part of it is a recommendation to buy or sell anything.

Trading Papers