Using the CtxLens Hierarchy

The CtxLens capture-hierarchy and how to use it.

1. Why this document exists

CtxLens is built around a hierarchy of context. The hierarchy shapes how cards are created, how they relate to each other, and how the eventual export reads. But the hierarchy is implicit in the product. It is not labeled on a screen, and a user can work in CtxLens for weeks without ever realizing it is there.

This document makes the hierarchy explicit, and it does so for a practical reason: a user who understands the hierarchy can capture far more granular information at the card level, with far less effort, than a user who does not. The difference is not subtle. Once the hierarchy is visible to you, your taxonomy gets sharper, your text entries get rarer and more specific, and the reports you produce get tighter and more useful.

This document is a working reference for the user who is already using CtxLens and wants to use it well.

How to read this document

Section 2 lists the full hierarchy in one place, end to end. Section 3 explains the practical insight that motivates the entire document (why the top two layers of the hierarchy are easy to overlook, and what gets lost when they are). Section 4 walks through every layer in turn, with examples drawn from the kind of work CtxLens is built for. Section 5 turns to the practical question of how to build the taxonomy that makes the hierarchy useful. Section 6 closes with a brief reflection on what the hierarchy produces over time.

2. The hierarchy in full

The CtxLens capture hierarchy, from highest (broadest) to lowest (most specific):

1. **Title** — set at export time in the Report Settings view.
2. **Subtitle** — set at export time in the Report Settings view.
3. **User** — implicit, attributed to the card's author.
4. **Timestamp** — implicit, set by the device at card creation.
5. **Group** — single-tap selection from your taxonomy.
6. **Subgroup** — single-tap selection, from your taxonomy, within the chosen group.
7. **GPS coordinates** — implicit if enabled in Settings; otherwise added per card.
8. **Tags** — multi-select, from your taxonomy, up to 24 per card.
9. **Images** — captured in-app or selected from your library, up to 12 per card.
10. **Text** — typed or dictated, up to 1,000 characters.

This ordering is not a prescription for the order in which you must populate elements as you capture. Order of entry does not matter. None of the non-implicit elements are programmatically required. What the hierarchy gives you is a framing: every characteristic you want to associate with an observation should first be attempted through an element other than text. Text is the last resort, not because text is unimportant, but because text carries the highest burden (on you writing it, and on the reader — human or machine). Everything above text in the hierarchy is either implicit, a single-tap selection, or a multi-select from your prepared taxonomy. Each is quick and easy to populate; each produces clean, structured context as a byproduct of how you already think about your work.

3. The structural insight

Notice that the top two levels of the hierarchy — Title and Subtitle — are not part of any individual card. You do not set them as you capture. You set them once, at the time you generate a report, in the Report Settings view. They sit above every card in the export.

This single fact changes how you might want to think about Group and Subgroup. Group is not the top of your reporting hierarchy. It is the third level down. You have two more layers above it that belong to the report itself, not to any individual capture.

If that sounds abstract, the practical consequence is anything but. It is the difference between burning two layers of organization on context that does not change all day and reserving those two layers for the actual structure of the work in front of you.

An example, with and without the hierarchy

Consider a user, call her Maria, who works at *Acme Engineering* in the *Mechanical Department*. Maria is investigating a pump failure on a particular asset, the *P-1203 centrifugal pump*.

A Maria who has not encountered the hierarchy might naturally do this:

- **Group:** Acme Engineering
- **Subgroup:** Mechanical Department

Her cards now sit under “Mechanical Department” inside “Acme Engineering.” She has used both of her granular organizational levels on information that is the same for every card she will capture for the next year. Anything she wants to say at the level of the actual investigation (which subsystem of the pump, which kind of inspection) has been pushed below where it belongs in the hierarchy. Some of it ends up squeezed into tags, where flat dimensions stand in for what should have been hierarchical structure. The rest ends up typed into the text field by hand, card after card.

A Maria who understands the hierarchy might do this instead:

- **Title** (set at export time): *Acme Engineering*
- **Subtitle** (set at export time): *Mechanical Department — P-1203 Pump Failure Investigation*
- **Group**: *Drive-End Bearing*
- **Subgroup**: *Race Inspection*

Her cards now organize themselves around the actual structure of her investigation. “Drive-End Bearing” and “Race Inspection” describe what each card is about. Her company and department still appear in the export; that information is contained at the top of the report, where it belongs, identifying the whole report rather than every card.

The first Maria has to type “DE bearing race showing spalling” into the text field of every relevant card, because Group and Subgroup are unavailable to her for that purpose. The second Maria has captured the same context in two taps and added a tag for “spalling” instead of typing the word fifteen times. Same investigation. Same evidence in the field. Two very different exports, and two very different experiences.

> *Title and Subtitle are set at export time. Group and Subgroup are the third and fourth layers down, not the first and second. Reserving them for the work itself is the single highest-leverage habit you can adopt.*

4. The layers, in turn

The remainder of this document walks through every layer of the hierarchy in order. Each section describes what the layer does, where it lives in the product, and how to use it well.

Title

The Title is the broadest level of context for an exported report. It is the first thing a reader sees. For most users, the Title identifies the work itself — the name of the project, the case number, the property address, or the company under whose name the report is being delivered.

The Title is set in the Report Settings view, immediately before you generate an export. It does not appear on individual cards. It sits at the top of the report, framing everything below it.

Subtitle

The Subtitle clarifies the Title. If the Title is “Acme Engineering,” the Subtitle might be “Mechanical Department — P-1203 Pump Failure Investigation.” If the Title is “1234 Main Street Appraisal,” the Subtitle might be “Subject Property and Three Comparables — Q2 2026.” The Subtitle sits directly below the Title in the export header. Like the Title, it is set in the Report Settings view, not on any individual card.

Together, Title and Subtitle establish what the report is. They are static and applicable across all cards in the card population selected for an exported report. They are the framing the reader needs before encountering any card. They are also where company name, department name, project name, and similar non-varying context belong: exactly the context that, in the absence of the hierarchy, users tend to put into Group and Subgroup and then regret.

User

Each card is attributed to its author. When you capture a card, you are the author. When a colleague transfers a card to you, the original author's attribution travels with the card.

Where authorship is displayed depends on the kind of report you generate and the options you select in Report Settings. In a standard report from a single author, your name typically appears in the export header, just below the Subtitle, because the entire report is yours and there is no need to repeat your name on every card. If you choose to include your name under the Subtitle, it appears there. If you generate a Report Optimized for use with External AI, your name appears explicitly per card, because an AI consuming the report benefits from knowing who authored each observation. If the export contains one or more transferred cards from a collaborator, every card carries its author's attribution — yours and theirs — because attribution is the basis of trust in a multi-author record.

The principle is simple: every card's author is permanently part of the record. CtxLens never strips authorship. It can adjust where authorship is displayed, but never whether it is preserved.

Timestamp

Each card is timestamped at creation. The timestamp is set by the device, automatically, with no input from you. It cannot be edited. It is not optional, though displaying it in reports is optional.

The timestamp is one of the most powerful elements in the hierarchy precisely because it requires no effort and produces high-quality structure. Cards in an export are ordered chronologically by default. A failure investigation reads as a story of when. A site walk reads as a route through time. None of that requires you to do anything beyond capture the cards as the work happens.

Group and Subgroup

Each card can be assigned to a group, and within that group, to a subgroup. Groups and subgroups are user-created and user-named. They are not predefined. You invent the taxonomy that fits your work, or adopt one from your company or your colleagues. Groups and subgroups can be created within CtxLens individually or built externally and added in bulk via a simple paste.

A mechanical engineer investigating a pump failure might have a group called *Drive-End Bearing* with subgroups *Race Inspection, Cage Inspection*, and *Element Inspection*. A commercial appraiser working a property might have groups for *Subject Property*, *Comparable Sales*, and *Neighborhood Context*, each with its own subgroups. The taxonomy is your mental model of your work, made visible in the product. The clarity of your mental model is what your reports will inherit.

Group and Subgroup are also re-nameable. In Settings, you can change what these layers are called across the entire app and in all exported reports. A construction inspector might rename Group to “Drawing” and Subgroup to “Detail.” A plant inspector might rename Group to “Unit” and Subgroup to “Equipment Type.” A surveyor will want something different again. You rename once to fit your profession, and the labels do their job throughout the user interface, on every card, in every export, and in every exchange with a collaborator. The purpose is fit, not fluidity.

A note on inheritance: Group and Subgroup carry over from card to card until you explicitly change them. No other element does. This saves you taps when you are working in a single area for a stretch of time, without introducing hidden errors — Group and Subgroup are prominent on the face of each card, so a stale carry-over is immediately visible. Tags, by contrast, would be easy to miss if stale, so they do not inherit. The exception is inside a single activity or countdown instance: later state-change cards in the same instance inherit Group, Subgroup, and tags from the starting card, because within one event those values are unlikely to vary and your attention is on the state change itself.

GPS coordinates

A card can carry a GPS location. If you enable it in Settings, GPS coordinates are added implicitly to each card as it is created. If not, you can add GPS coordinates to any card at any time, and leave them off others. The choice is yours, and it is per card.

GPS becomes especially valuable when a single body of work spans multiple physical locations — a commercial appraiser comparing a subject property against three comps across a neighborhood, an architect documenting sun angles from different vantage points on a site, a forensic engineer tracking the distribution of failures across a plant. In those cases, the GPS on each card is the spine of the spatial story the export tells.

For cards captured at a desk, GPS is likely unnecessary noise. You decide, per card, whether the location was meaningful.

Tags

Each card can carry up to 24 tags. Tags are named and created by you. Like groups and subgroups, they can be added to CtxLens individually or built externally and bulk-pasted into CtxLens. Unlike groups and subgroups — which place a card in a hierarchy — tags describe a card along orthogonal dimensions. A single card can have many tags simultaneously.

Tags are where granularity is established at the lowest cost to you. A well-built tag taxonomy lets you tag a card in seconds in ways that make it searchable, sortable, and contextually precise later. A card might carry *evidence:direct-observation, confidence:high, mode:fatigue, component:DE-bearing,* and *relevance:rca-primary*. Five tags. Four seconds of tapping. The card now carries enough structure that a downstream consumer of the export (another human or an AI) can situate it within the whole body of work without needing to re-read the free text.

This is the element of CtxLens that most rewards deliberate taxonomy work. A taxonomy specific to your work is easy to create and extremely valuable to both capture and reporting. There are several ways to develop one: build it in advance, develop it iteratively from the cards themselves, or use AI to surface it from your captured work. All three paths are treated in the section that follows. In every case, your tag taxonomy is your context-engineering language. It is the vocabulary your exports will speak in.

A note on display order

In Settings, you can choose how Group, Subgroup, and Tags are ordered when you tap the picker for them: alphabetically (ascending) or by the order in which you created them. Pick whichever matches how you reach for entries. Alphabetical works well for taxonomies with many entries, where you scan by name. Order-created works well when you would like to see your Groups, Subgroups and/or Tags in a specific order (create them in that order and they will be displayed throughout the app in that order). The setting is one tap to change, so trying both and seeing which feels faster is the simplest way to choose. For users who want to control the picker order more deliberately than either default allows, see Appendix A.

Images

Each card can carry up to 12 images. You can capture them fresh with the in-app camera, or select them from your photo library. Both paths are supported because both are needed: the camera is for the moment, the library is for reference material you already have. A card might contain three photos taken in the field and one screenshot from an applicable regulatory document. All four belong in the same card because they support the same observation.

CtxLens does not enhance, transform, or auto-caption your images. They are yours, as taken. If you want them annotated, you may annotate them before adding them from your library.

Text

Each card has one free-text field, up to 1,000 characters. The field is where you write the judgment that the other elements cannot express. “Bearing temperature reading 87°C, 12°C above baseline, climbing over last 30 minutes per operator report” is a card text. “Homeowner prefers warm wood tones — specifically called out white oak on cabinet sample C-7 during kitchen walkthrough” is a card text. The text is a captured observation, spoken plainly, by the person who was there. It picks up where the structured elements left off.

Text can be entered by typing or by dictation. Dictation uses the iOS keyboard's built-in microphone; the result lands in the field as plain text, capped at 1,000 characters like any other text entry. The cap applies to dictated text for the same reason it applies to typed text: the field is for a specific residual observation, not for a long narrative.

If you have read the rest of this document, the role of text is clear. Text is the last resort, not because text is unimportant, but because text is the most expensive element to produce and the least structured. Everything above text in the hierarchy is either implicit, a single-tap selection, or a multi-select from a prepared taxonomy. By the time you reach for the text field (if you reach for it at all) the context around your observation is already established. Title, Subtitle, User, and Timestamp have set the framing. Group and Subgroup have set the domain. Tags have set the dimensions. Images have set what you are looking at. GPS has set where this is taking place. You can now speak directly from inside that established context, rather than having to write the scene before making your point.

For a well-prepared user, the residual observation that text is for rarely requires more than a hundred characters.

Transferred cards

Cards can be received from a collaborator, without a server, without a network connection, and without administrator setup. A user working on a body of work can receive cards captured by a trusted colleague and incorporate them chronologically with their own. This matters for field work spanning more than a single person — a joint site visit, a two-engineer investigation, an architect working with a surveyor — because the alternative is to re-enter the colleague's work, which defeats the premise.

The rules around transferred cards are precise.

The authoring user's content is locked. A transferred card's GPS, text, images, group, and subgroup cannot be modified by the receiving user. Whatever the original author captured is preserved exactly as they captured it.

The receiving user can, however, add their own tags. This matters because the receiving user may carry context the original author did not have. Adding tags allows that cumulative context to attach to the card without overwriting the authoring user's work. Every tag added by the receiving user is attributed to them in the database, in the user interface, and in every export format.

When transferred cards are present in an export, every card carries its author's name on its face, because the report becomes a multi-author record and authorship needs to be visible per card. This is the case whether or not you have selected to include your own name under the Subtitle. The presence of a collaborator's work changes what the export needs to communicate.

The export

Your body of work — your title and subtitle, your cards, your groups, your subgroups, your tags, your GPS, your images, your transferred cards — can be exported in a number of formats, each suited to a different audience and purpose. The full menu of export options, and what each is designed for, is described in *CtxLens Reports*, a separate technical document. The Report Settings view inside CtxLens is a good place to explore the formats directly.

The Report Settings view is also where you populate the top two layers of the hierarchy. It is where you set the Title, the Subtitle, the type of report, and whether your name should appear under the Subtitle in the export header. It is the moment where the framing of the report is finalized.

5. Building your taxonomy

Knowing what each layer of the hierarchy does is one half of using CtxLens well. The other half is having a taxonomy that represents your work. Your tags, your groups, your subgroups — these are not given to you by the product. You build them. The question is how.

There are three paths to a working taxonomy. The first is appropriate when you already know what you need. The second is appropriate when you do not. The third uses the cards you have already captured, with AI as a pattern-recognition assistant, to surface a taxonomy fitted to your actual work. All three produce strong results. Most users end up using a blend.

The bulk path: when the taxonomy is known

Some users come to CtxLens with a taxonomy already in hand. Common sources include:

- An existing standard from your employer.
- A recognized classification system in your industry.
- A colleague's tag list you have adopted.
- Years of accumulated experience that gives you a clear vocabulary for the work you do.

When the taxonomy is known, the bulk path is the right one. CtxLens accepts the bulk creation of groups, subgroups, and tags via simple paste. You can prepare the list externally (in a spreadsheet, a text document, even a chat message) and bring it in once. After that, every card you capture has the full taxonomy available with a tap.

For users whose work is governed by a recognized standard, this path can take as little as a few minutes and pay back across years of card creation in CtxLens. For users joining a team that already uses CtxLens, the bulk path is how the team's vocabulary becomes yours. The cost is upfront and small. The benefit accrues forever.

The iterative path: when the taxonomy is not yet known

Most users do not arrive with a complete taxonomy. They know roughly what their work involves, but they do not yet know which observations will repeat, which dimensions will matter most, or which words their future self will reach for to describe a recurring pattern. That is fine. CtxLens is designed to let the taxonomy emerge from the work itself.

Here is how this works in practice. You start capturing cards. Group and Subgroup get populated naturally — even an early-stage user usually knows what large bucket their work falls under. Tags may be sparse at first. The text field does more of the heavy lifting than it will later. None of this is a sign that you are using the product wrong. It is the starting state of the iterative path.

After a session — at the end of a day, the end of a site visit, or the end of a project phase — you review the cards you have created. Read the text field of each one. Some of those text entries will be specific to a single card and will never recur. Leave them alone; the text field is doing real work there.

But others will repeat. You will see the same descriptor coming up in three, five, ten cards. “Spalling visible.” “Within manufacturer's tolerance.” “Not addressed in original specification.” “Moisture damage in framing.” “Witness reports indicate.” These repeating descriptors are the signal. They are observations that are not residual. They belong in the taxonomy.

For each repeating descriptor, create a tag — or, if the descriptor is broader than a single tag, a subgroup or a new dimension in your tag taxonomy. Use it on future cards. The text field for those cards will get shorter, because the descriptor that was previously consuming words is now a one-tap selection. Capture will get faster. And the granularity of your eventual report will deepen, because what was previously buried inside free text is now visible as structured metadata that can be filtered, counted, and reasoned over.

The AI-assisted path: surfacing your taxonomy from existing cards

There is a third path worth describing, and it sits between the other two. It is most useful once you have a meaningful corpus of cards (fifty, a hundred, several hundred) that you have captured iteratively over a period of work. By that point, your text fields will likely contain many recurring descriptors and reviewing them card by card to identify the repetitions is tedious. AI is good at exactly this kind of pattern recognition. CtxLens is built to take advantage of that.

The workflow is straightforward. Select a large population of your existing cards and generate an export using the Report Optimized for use with External AI format. That export is structured for AI consumption — it carries explicit per-card authorship, the full text of every card, the full element set, associated images, and in a layout downstream models can parse cleanly. Attach the resulting PDF to your preferred AI and ask it explicitly to review the text fields across all cards, identify concepts that recur, and propose additional tags to add to your taxonomy that would capture them. A prompt as plain as “Review the text field of every card in this report. Identify concepts and descriptors that recur across multiple cards. Propose a tag taxonomy in dimension:value form, one tag per line, ready to copy and paste” tends to produce a usable result.

The AI's proposed taxonomy can then be reviewed, edited, copied, and pasted into CtxLens's Bulk Add field for tags. The same approach works for subgroups, particularly when the subgroup field has been underused or used inconsistently and the AI's review of card text reveals a clearer organizing structure than the one currently in place.

It is worth being precise about what this workflow is and is not. The AI is not generating context for you. It is identifying patterns in the text you already authored, with your judgment, at the moment of capture. The taxonomy it surfaces is your taxonomy. The AI is doing the tedious recognition work (finding the descriptors that repeat) that you would otherwise have to do by hand. You retain the judgment about which suggestions to accept, which to edit, and which to discard. AI proposes; you decide.

Once the bulk-paste is complete, you are back in the same position as a user who arrived with a known taxonomy in hand — except this taxonomy was extracted from your actual work, which means it fits that work more precisely than any external standard ever could. From here, the iterative path resumes for any descriptors that emerge later. The feedback loop continues exactly as before.

The feedback loop

The paths converge. After enough iterations, your taxonomy reaches a point where new cards rarely require anything in the text field beyond a sentence or two of genuinely specific observation. The repetitive language has all been promoted to tags. The structural language has all been captured by groups and subgroups. The text field is reserved for what only you can say about what you saw.

This is the feedback loop the hierarchy is designed to produce. Every time you notice yourself typing the same words across two consecutive cards, you have an opportunity to make capture cheaper going forward. Take it. Add the tag. The cost is a few seconds of taxonomy work. The payoff is an indefinite reduction in typing across every card that comes after, and a corresponding gain in the granularity of your reports.

Seen this way, the three paths are not really separate. They are the same work (extracting a taxonomy from a body of expertise) performed at different moments and with different assistance. A user who builds a complete taxonomy in advance has done the work mentally, before opening CtxLens, drawing on prior experience. A user who builds it iteratively is doing the same work, just visibly and in real time. An AI-assisted user is delegating the recognition step while keeping the judgment step. All three end up in the same place: a taxonomy that fits the work, populated with confidence, producing strong exports.

A note on naming conventions

Tags that follow a consistent naming pattern are easier to use, easier to find, and easier for AI to reason about. A common convention is *dimension:value*, as shown earlier in the Tags section — *evidence:direct-observation, confidence:high, mode:fatigue*. The dimension before the colon tells you what kind of fact the tag is recording. The value after the colon is the specific fact. This pattern keeps tags grouped intuitively in the tag selector, and it makes the resulting export easy for any reader, human or otherwise, to parse.

There is no enforcement of any naming convention in CtxLens. You can use any pattern you like, or no pattern at all. But once you have settled on one, stick with it. A taxonomy with two patterns is harder to use than a taxonomy with one.

6. Putting it to work

The hierarchy is not a rule you have to memorize. After a few sessions with the hierarchy in mind, it becomes natural. You will find yourself reaching for tags before text. You will find your taxonomy growing a little more specific each week. You will find your text entries getting shorter and more precise.

Note-taking apps make writing the rule and structure the exception. CtxLens makes structure the rule and writing the exception. The hierarchy is the mechanism that makes the exception possible. Title and Subtitle frame the report at export time. Group and Subgroup organize the work itself. Tags add dimensional precision. Images and GPS fill in the where and the what. Text is reserved for the small residual that nothing else can carry. Each layer does its job. Each layer is cheaper to produce than the layer below it. The cards you create begin to carry a richness that typing alone would struggle to produce.

That is the goal. CtxLens is built to make capture efficient enough to become a habit. The hierarchy is what makes the habit produce something worth having.

Appendix A: Controlling display order in the picker

CtxLens lets you choose how Group, Subgroup, Tags (Actvities and Countdowns) display in the picker — alphabetically or by the order in which you created them. That setting is described briefly in Section 4. Most users pick one and move on.

For users who want to control the exact order in which their taxonomy values appear during capture — not just alphabetical, not just chronological, but a specific order that mirrors how they think about the work — there is a more deliberate move available, and it is worth knowing about.

When CtxLens commits a bulk-pasted set of values to its database, it explicitly separates the database commit of each item by one millisecond. The reason is to ensure that when you have chosen “order created” as your picker sort (in Settings), a bulk-imported list still respects the order of the paste. Without that separation, every item in the bulk paste would share an identical timestamp and the picker would render them in an arbitrary internal order. The 1-millisecond stagger guarantees deterministic, paste-order rendering.

The implication for power users is meaningful. If you have built your taxonomy iteratively over a period of time and you want to control the exact order in which your values appear in the picker, and throughout CtxLens, you can:

1. Export your existing taxonomy values to text.
2. Reorder the list externally, in whatever order you want it to appear.
3. Delete the taxonomy values inside CtxLens.
4. Bulk-paste the reordered list back in.

With “order created” selected as your sort, your picker will now display the values in exactly the order you chose. Any new values added later will append to the end — or, if you redo the rebuild, to wherever you place them in the next paste.

This workflow is overkill for users with small or alphabetically-organized taxonomies. It is genuinely useful for users with large taxonomies whose mental model has a natural ordering that is neither alphabetical nor chronological — for example, an inspector whose tags should mirror the natural sequence of an inspection, or an appraiser whose subgroups should mirror the order of a property walk. For those users, the rebuild workflow is the difference between a picker that fights you and a picker that mirrors how you think.