



A FREE RESOURCE FROM

The Genealogist's Prompt Library

Ready-to-use AI prompts for careful, evidence-based family history research

Presented by **Norwegian Genealogy and Then Some**
martinroe.com/blog

The prompts in this library are meant to be reused. They are written as working instructions rather than as questions with a single correct answer. The bracketed sections are where you would put the details of your own research problem.

A prompt becomes more useful when the AI knows what you already have and what you want it to do with that material. That distinction runs through the whole collection. If you want a research strategy, ask for a strategy. If you want a transcription, ask for a transcription. If you want a translation, give it a checked transcription and ask it to translate that text. The wording matters, because an AI system will otherwise tend to fill the space with an answer of its own.

Use these prompts as starting points. They are deliberately plain. There is no advantage in making a prompt sound clever. A good prompt tells the AI what information it may use, what task you want performed, and where you do not want it to go beyond the supplied material.

PLANNING YOUR RESEARCH

Research Strategy

When you have a person whose life you are trying to reconstruct, give the AI the evidence already collected and ask it to build a research strategy.

"I am researching [PERSON]. Here is what I know: [EVIDENCE]. My research question is: [QUESTION]. Develop a research strategy rather than an answer. Prioritise the next searches by how directly they could answer the question. Tell me what each search could establish. Do not invent records or sources."

The last sentence keeps the task tied to the material and to searches that can actually be carried out. This asks for a route through the research, not for the missing ancestor to be produced on demand. Put the research question in its own sentence rather than burying it in a long description — "Where was this person living in the period between the two known records?" gives the AI a defined problem to work with.

Research Brief

When notes have grown into a mixture of copied information, comments to yourself, and unanswered questions, this prompt turns that material into something easier to work with.

"Turn these notes into a concise research brief. Separate confirmed facts from assumptions and open questions. Preserve source information. Do not add facts."

Preserving source information matters — if a note says where a fact came from, that should stay attached to it. Use this before beginning a new search: a concise brief is a cleaner starting point than several pages of old notes, and it shows which parts are established and which are still your own working assumptions.

Next Record

When you know your current research question but are unsure which record to look for next.

"Based on the confirmed evidence below, identify the record that could most directly answer my current research question. Explain what information I should look for in that record."

Supply the confirmed evidence rather than a collection of old suggestions — that keeps the distinction clear between what you have actually established and what you have merely considered.

Negative Search

A search that produces no result can be difficult to interpret.

"I searched [SOURCE] for [PERSON/PERIOD/PLACE] and found no matching record. Explain what this negative result establishes and what it does not establish. Suggest the next search without assuming that the record does not exist."

A failed search tells you something about that search — it does not automatically tell you the person or record never

existed. The next search can be directed at the limitations of the first one.

What Am I Missing?

Sometimes the problem is easier to describe as a missing piece than as a theory.

"Review the research below. Identify one important piece of evidence or one research question that is still missing. Explain why it could change the current assessment. Do not invent facts."

Asking for one missing piece keeps the response focused — this is about identifying a gap in the material supplied, not producing a list of every record that might exist. The "why it matters" part helps decide whether a search is worth pursuing or merely an interesting possibility.

TESTING AN IDENTIFICATION

Same-Name Problem

Two people with the same name can easily become one person in a genealogy database.

"I have two people with the same name. Compare them using the supplied age, residence, occupation, relatives and dates. Identify evidence that distinguishes them. Do not merge the identities because the names match."

A name by itself is not enough for this task — age, residence, occupation, relatives, and dates provide real points of comparison when those details are present in the records. The final instruction prevents the name itself from becoming the conclusion.

Hypothesis Test

Useful when you have already formed an idea about what happened.

"My hypothesis is [HYPOTHESIS]. The evidence is [EVIDENCE]. Test the hypothesis. Identify evidence that supports it, evidence that weakens it, and information that is missing. Do not assume the hypothesis is correct."

This asks the AI to work with a hypothesis you supply, rather than inventing one, and to identify missing information rather than filling the gap itself. Use it when you've reached a possible conclusion and want to examine it before treating it as established — the prompt does not ask the AI to prove you right.

Devil's Advocate

Takes a conclusion you currently hold and asks the AI to challenge it.

"I currently believe [CONCLUSION]. Use only the evidence below. Look for reasons the conclusion could be wrong. Identify the strongest alternative explanation and the evidence that would distinguish it."

The request for an alternative explanation, and for the evidence that would distinguish it, gives you something concrete to look for in the records. Use this after a promising line of research has started to feel settled — a conclusion can become familiar simply because you have repeated it in your notes.

Emigration

For a person who appears in a Norwegian record and later in another country.

"I have a Norwegian record for [PERSON] and a later record in [COUNTRY]. Compare the supplied details and identify evidence that could connect the identities. Do not assume that they are the same person."

The instruction against assuming identity keeps the two records separate until the evidence actually connects them.

Census Extraction

For pulling structured information out of a census entry.

"Extract names, ages, relationships, occupations and residences from this census entry. Preserve the recorded information. Do not infer relationships that are not stated."

This is an extraction task, not an interpretation task — important when a household includes people whose relationship to each other isn't explicitly stated. The extraction should reflect the entry, not an assumed relationship.

Census Comparison

For comparing two census households.

"Compare these two census households. Identify people who may correspond across the records and explain the evidence for each proposed correspondence. Mark uncertain identities."

The task is comparison, not automatic identification — you get a working comparison you can check against the records themselves.

Church Record Transcription

Handwritten church records create their own problem.

"Transcribe this church-record entry as closely as possible. Preserve names and place names. Mark uncertain readings. Do not guess unclear handwriting from the surrounding family."

The final sentence is deliberately strict. Familiarity with a family can tempt anyone reading a difficult entry to supply the name they expect to see. A transcription that preserves uncertainty is more useful than one that silently replaces an unclear word with a plausible one.

Church Record Extraction

Once you have a checked transcription.

"Using the checked transcription, extract every genealogically relevant name, date, place and relationship. Do not add information."

This is a different job from transcription — you're no longer asking the AI to interpret handwriting, only to organize information in text you've already checked.

Passenger Record

For a passenger-list entry.

"Extract the identifying information from this passenger-list entry. Preserve the recorded spelling, age, occupation, residence and destination. Do not add information."

Preserving the recorded spelling is especially useful when a name has been altered elsewhere in your notes.

Translation

Translation comes after transcription, not instead of it.

"Translate this checked transcription into English. Preserve personal and place names. Explain historical terms that do not have a direct modern equivalent."

This assumes the transcription is already checked. Where a historical term has no direct modern equivalent, the AI can explain it rather than forcing a misleading translation.

NAMES AND PLACES

Farm Name

Norwegian records can contain both a patronymic and a farm name.

"The person appears with [PATRONYMIC] in one record and [FARM NAME] in another. Treat the farm name as a place unless the evidence shows otherwise. Explain what evidence could establish how the name is being used."

This tells the AI to start from the farm name as a place — it doesn't get to become a surname simply because it follows a given name in another record. The question of how the name is being used stays open until the records provide the answer.

Place-Name Question

A place name can require the same restraint as a personal name.

"The record contains the place name [PLACE]. Identify what information I need before assuming which location it refers to. Do not choose a location without evidence."

This asks for the information needed to distinguish between possible locations, rather than asking the AI to pick one because it seems likely.

RESOLVING CONFLICTS

Conflicting Dates

Records sometimes give different dates for the same event.

"These records give different dates for the same event. Compare the dates without choosing one. Identify possible explanations and the evidence that could resolve the conflict."

Useful when you haven't yet found enough evidence to decide between them — the AI can identify possible explanations and point toward the kind of evidence that might settle the disagreement.

Conflicting Residences

The same approach applies when residences differ.

"These records give different residences. Identify what each record actually establishes. Do not invent a move date. Suggest records that could document the change."

A different residence in two records does not, by itself, give you a date for a move. This keeps the two facts separate and asks what records could document the change.

Source-Bound Analysis

For keeping the AI entirely within a supplied set of records.

"Use only the records supplied below. Do not use earlier AI suggestions as evidence. If the evidence is insufficient, say so."

Especially useful in a long conversation, where earlier suggestions can linger even after later research has changed the picture. This tells the AI which material should count for the current task.

Source Critic

For checking a proposed conclusion against the sources.

"Review this proposed conclusion against the supplied sources. Identify every statement that is directly supported, inferred or unsupported."

This lets you examine a paragraph or conclusion statement by statement, separating direct support from inference rather than treating the whole thing as one unit.

DOCUMENTING YOUR WORK

Family Group

For turning a set of records into a family-group summary.

"Create a family-group summary from these records. Separate information directly stated in a source from relationships inferred by comparing sources."

The distinction between stated information and inferred relationships is easy to lose when several records are read together. Use this with records you've already checked, rather than asking the AI to supply missing family members.

Timeline

Straightforward when the records contain clear dates.

"Create a chronological timeline from the supplied records. Include date, place, person and source. Do not create dates for undated events. Mark estimates as estimates."

If a record has no date, you don't want one appearing in the timeline simply because the surrounding entries have dates. An estimate should always look visibly different from a recorded date.

Research Log

For turning working notes into a consistent research record.

"Turn these notes into a research-log entry containing the research question, search performed, source consulted, result and next step. Do not add facts."

The fixed fields give the log a consistent structure; the final instruction keeps the AI from filling gaps in your notes.

End of Session

For the end of a research session.

"Based only on the information in this conversation, separate established facts, inferences, hypotheses and unanswered questions. Do not introduce new information."

Creates a clean division between what's established and what remains open — useful before closing a session, particularly when the conversation has run long.

Clean Reset

When a conversation has accumulated old suggestions.

"Use the following as the current evidence set. Ignore earlier assumptions and AI suggestions that are not included here."

A simple reset — place the material you want to work with after the instruction, and leave earlier material outside the

current evidence set.

Update After New Evidence

New evidence can change an earlier assessment.

"Here is new evidence. Reassess the current hypotheses. Identify what has changed and which earlier conclusions need to be reconsidered. Do not defend an earlier answer merely because it appeared earlier in the conversation."

The final sentence addresses a real problem in long AI conversations — an earlier answer can stick around simply because it was written first. The new evidence should be assessed on its own terms.

Final Assessment

When you've finished a stage of the research.

"Write a research assessment using only the checked evidence. Separate facts from inference. State what is currently supported and what remains unresolved. Do not invent citations or facts."

Deliberately tied to checked evidence only — the AI states what the material supports and what remains unresolved, without adding anything from elsewhere.

USING THIS LIBRARY

These prompts are starting points, not fixed scripts. Add enough context for the AI to understand your specific problem, and remove any instruction that doesn't apply — that principle matters more than the exact wording.

A prompt can be short when the task is simple. A complicated research problem usually needs more context. If you're asking about two people with the same name, provide the details that let them be compared. If you're asking about conflicting residences, provide the records containing those residences. If you're asking for a translation, use the checked transcription rather than the original handwriting.

Avoid carrying unnecessary instructions from one task into another. A transcription prompt and a translation prompt have different jobs — the first deals with what is written in the record, the second with the meaning of a checked transcription. Combining every possible instruction into one enormous prompt makes it harder to see what you're actually asking the AI to do.

The prompts can also be combined when the research requires it. You might first create a research brief from your notes, then ask for the next record, and later build a research log from the searches you performed. A census entry might be extracted, compared with another household, and then folded into a family-group summary.

Think of this as a set of tools kept within reach. You don't need every tool for every problem. When you're staring at a difficult church-record entry, you need the transcription prompt. When two census entries are sitting side by side, you need the comparison prompt. When your notes have become tangled, you need the research-brief prompt.

These prompts also give you a useful discipline in how you work with AI. You have to decide what you're asking it to do before you ask — that sounds obvious, but it's easy to type a long description of an ancestor and end with "What do you think?" The resulting answer may contain useful ideas, but the task itself was never well defined.

A more specific instruction gives you something you can check. You can ask the AI to extract information from a census entry and compare its extraction with the entry itself. You can ask it to identify a missing research question and decide whether that question is worth pursuing. You can ask it to review a proposed conclusion and inspect the statements it marks as inferred.

Keep the prompts themselves somewhere easy to reach. The point of a prompt library is reuse — there's little value in a good prompt if you have to reconstruct it from memory every time you meet the same research problem. Expect to alter the wording as you go: one problem may need more context than the brackets suggest, another may need less. The rule that matters most is simple — add enough context for the AI to understand the particular problem, and remove any instruction that doesn't apply.

That is how to use a prompt library: choose the task first, supply the material needed for it, and tell the AI where the boundaries are. The resulting answer is then something you can inspect against the records you actually have.