

Recreating property boundaries from old maps and cadastral records using AI: The case of Røst, Norway

Sjur K. DYRKOLBOTN, Arve LEIKNES and Leiv Bjarte MJØS, Norway

Key words: AI, cadastral surveying, LLMs, agents, geospatial reasoning

SUMMARY

Property boundaries in rural Norway are often recorded only in historical legal texts and sketch maps rather than as georeferenced coordinates. This paper presents a case study from the island of Røst, where a recent formal cadastral survey provides partial ground truth for evaluating AI-generated boundary proposals.

We study the use of multimodal large language models (LLMs), enhanced by an approach we call the puzzle strategy: instead of processing each parcel in isolation, we first prompt the model to reason holistically about how all parcels in a cluster fit together — spatially, historically, and legally — before asking for coordinates. This mirrors the structure of legal argumentation under conflicting evidence: inconsistencies are treated as conflicts to be resolved by appeal to a priority ordering over sources, rather than patched by an arbitrary geometric closure.

The preliminary results show a clear improvement over direct, parcel-by-parcel elicitation. We argue that structured relational prompting of this kind could form the basis of a practical AI copilot for cadastral surveyors, with applicability beyond Norway wherever historical records are the primary source of boundary information.

Recreating property boundaries from old maps and cadastral records using AI: The case of Røst, Norway

Sjur K. DYRKOLBOTN, Arve LEIKNES and Leiv Bjarte MJØS, Norway

Key words: AI, cadastral surveying, LLMs, agents, geospatial reasoning

1. INTRODUCTION

In many countries, one may encounter property parcels that are properly documented only in historical documents rather than in modern spatial databases. Norway is a striking example: a rich archive of historical land-division records exists, but translating this archive into georeferenced geometry has traditionally required time-consuming and expensive expert assessment, followed by field work. No systematic high-quality surveying of already existing properties has ever taken place. A modern cadastre has nevertheless been introduced, with seemingly precise boundary information provided for most parcels. Much of this information is based on inaccurate sources, however, including imprecise economic mapping using orthophoto-based techniques. Hence, information about boundaries in the cadastre often *conflicts* with the historical documents on record, as well as the physical boundaries in the terrain, resulting in many boundary disputes (Leiknes and Mjøs, 2021). It is therefore important to think about how to improve the quality of the information in the cadastre in a systematic way. Recent advances in artificial intelligence (AI), particularly the emergence of highly capable large language models (LLMs), suggests a possible path forward. In this paper we report some preliminary results on this possibility, including a relational approach to LLM-assisted surveying that we call the *puzzle* strategy.

2. BACKGROUND

Until 1980, the *skylddeling* (“dividing the duty”) was the standard instrument for subdividing rural properties in Norway. This reflected a value-based approach to property registration, focused on recording the relative value of property parcels (originally for rent or tax purposes). The resulting official document when a property was subdivided, a *skyldskifte*, also called a *skylddeling*, describes boundary points and the directions of boundary lines, normally also the boundary lengths. It does so textually — often by reference to natural features such as streams, boulders and roads, and to neighbouring estates — but rarely with maps or coordinates. When working in rural areas, these documents are widely used by surveyors, first when looking for boundary marks (which tend to be decisive when found), then to fill gaps left by marks that are missing or disputed. Other important sources of information in this

setting are the so-called *utskiftninger* (land-consolidation maps and proceedings), mostly from the 19th and early 20th centuries. These documents describe comprehensive restructurings of farms that had previously been subdivided in a fragmented manner unsuitable for modern agriculture. The typical consolidation case involves several properties, often adding hundreds of pages of information per case, usually also with detailed annotated maps. A consolidation would commonly result in a new start for the local property structure that then served as a basis for further subdivision by *skyldskifte*. A good approach to surveying is therefore to start from the first consolidation and work chronologically up to the present day, tracking adjustments and subdivisions along the way. This is a time-consuming process that often depends on expert knowledge and painstaking cross-referencing across many different documents, as well as costly field work.

Unfortunately, a systematic approach is the only approach that is likely to work if the aim is to significantly improve the cadastre within a reasonable timeframe. If errors are only dealt with on a case-by-case basis, as is normal in Norway, new property formations and adjustments will tend to proliferate errors more quickly than they are resolved, making the problem *worse* over time. The underlying issue is that a correct cadastral map is not merely a collection of polygons that seem correct in isolation, but a consistent tiling of a large area by property parcels that all fit together correctly. Boundaries are inherently *relational*.

This is true in general, suggesting that a holistic approach to property registration is also far more *efficient* in the long run, if it is properly funded and organised (Larsson, 1991, pp. 81-82). This becomes particularly clear when working with old documents that lack precise coordinates: a description such as "runs northwards along the western boundary of the farm to the road" has limited geometric content in isolation. Any reconstruction, even for a single contested or uncertain boundary, should reason about the entire neighbourhood, but will normally not do so. When applied to properties formed by *skylddeling*, the parcel-by-parcel approach traditionally adopted in Norway may well solve one problem while creating several new ones. A more holistic approach is therefore needed. Moreover, one must consider the parcel-formation history: most parcels have a documented formation event, and a child parcel must usually be geometrically contained within its parent at the time of subdivision.

Recent advances in multimodal large language models open new possibilities for cost-effective systematic surveying. Such models can already read historical Norwegian texts, extract topological boundary descriptions, reason about spatial relationships, and generate approximate coordinate sequences that describe plausible property parcels. For textual deed analysis, LLMs have already shown considerable promise, and commercial products are starting to appear.¹ However, in our experience, they process one document at a time and fail

¹ See, e.g., https://buntinglabs.com/solutions/digitize-survey-plot-pdf?trk=public_post_comment-text and <https://www.v7labs.com/agents/deed-analysis-agent>.

to reason effectively about relational, incomplete and ambiguous information, making them unsuitable for Norwegian historical records.

In the following, we present a case study that illustrates both the limits of the parcel-by-parcel approach as well as the potential of a more holistic method. The key idea is to use the LLM to process large amounts of data and help with the spatial and relational reasoning.

3. CASE STUDY: GNR. 12, NEDRE NES, RØST

3.1 Study area

Røst is a small island municipality at the south-western tip of the Lofoten archipelago in Nordland county, Norway (Figure 1). With approximately 500 inhabitants and an economy historically based on fishing, Røst is a typical case of coastal Norway: a web of small parcels formed through repeated historical subdivisions, with boundaries documented in a mixture of *skyldskifter* and *utskiftninger*. A recent cadastral survey of part of the island, conducted by staff and students from the Western Norway University of Applied Sciences (HVL), provides partial ground truth for evaluating AI-generated boundary proposals.



Figure 1. *Aerial photograph of the Nedre Nes study area, Røst. The cluster of buildings at the centre corresponds to the properties in Gnr. 12. The rocky shoreline is visible to the north.*

3.2 The parcel cluster and available data

The focus area is the cadastral cluster at Gnr. 12, Nedre Nes (Figure 2). Before the recent survey, several parcels in this cluster had boundaries recorded in the cadastral map that were inconsistent with the documentary evidence and the boundary marks. The data used in the testing comprise:

- The *utskiftning* for Nedre Nes, gnr. 12.
- *Skyldskifter* covering parcels 12/20, 12/21, 12/22, 12/30 and 12/38, describing boundaries by reference to roads, fences and neighbouring parcels.
- An aerial photograph providing visual reference for identifiable features (roads, fences, coastline).

The *skyldskifter* contain internal inconsistencies: distances and bearings from different documents cannot always be reconciled, and polygons don't close properly with the information provided (assuming measurements are accurate). These characteristics make the cluster an ideal case for testing how LLMs can help us deal with typical inaccuracies and inconsistencies encountered when working with these sources. Moreover, we only provided the model with a subset of all the available documents, leaving intentional data gaps. This was done to test how LLMs can help us deal with incomplete information, which is also typical.



Figure 2. Extract from the cadastral map (*matrikkelen*) showing the Gnr. 12 / Nedre Nes study area. GPT-5 was asked to recreate the boundaries of 12/38, 12/20 and 12/21 using only *skyldskifter*, an *utskiftning* and an aerial photograph.

4. EXPERIMENTS AND RESULTS

We conducted three experiments with GPT-5, each varying the amount and structure of information provided to the model.

4.1 Experiment 1: Direct geometry elicitation

In the first experiment, GPT-5 was given only the textual boundary description for a single parcel (12/38) and the coordinates of an anchor point at the south-east corner. It was asked to produce a georeferenced polygon. The model correctly identified the anchor in the text and applied the described distances and bearings to construct the parcel, returning valid GeoJSON. This baseline mirrors the parcel-by-parcel workflow of current commercial tools.

The output was qualitatively reasonable but geometrically imprecise, and three systematic problems emerged:

- *Source inconsistency and non-closure.* The *skyldskifte* for parcel 12/38 was internally inconsistent: reported distances and bearings did not allow the polygon to close. GPT-5 correctly noted this, but in the geometric output it resulted in an anomalous southward protrusion — an implicit conflict-resolution for an inconsistency that should instead have been flagged for explicit consideration.
- *Geometric displacement.* With only one anchor and unresolved inconsistencies the constructed polygon was displaced relative to the ground truth.
- *Attempted repair.* Without additional evidence, the model had no principled basis for choosing how to resolve the inconsistencies; the resulting geometry was only partly determined and the model made arbitrary choices to complete the task. The protrusion was also an *unreasonable* choice that no human would have made.



Figure 3. Result of Experiment 1 (direct elicitation). GPT-5 polygon for parcel 12/38 shown in blue; cadastral map boundary in red. The protrusion to the southeast arises from the model's attempt to close the polygon without violating the measurements and bearings given in the *skyldskifte*.

4.2 Experiment 2: Relational reasoning and attention steering

In the second experiment, the *skyldskifter* for all parcels in the cluster were combined in a single context. Rather than asking for a polygon immediately, we first asked GPT-5 to reason in natural language about the spatial structure: which parcels share boundary segments; how the road running south of the cluster constrains the southern boundaries; what the parcel-formation history implies about containment; and where internal inconsistencies lie and how they might be reconciled. Only after a detailed relational account was produced did we ask for coordinates of three parcels simultaneously, anchored on a single known road intersection to the south-east. A georeferenced aerial photograph was provided as an additional reference point for the model, showing the road and thereby allowing the LLM to infer a better interpretation of the imprecise bearings provided in the texts.

The explicit relational reasoning step served as an attention-steering mechanism, helping address the “lost in the middle” phenomenon addressed in (Liu et al. 2024), keeping the model focused on the data points most likely to be relevant to the main task, despite providing it with a long context. Effectively, we prompted the model to articulate cross-document constraints and hold them in working context before generating geometry.

When the relational reasoning approach was applied to all *skyltskifter* in the cluster simultaneously, the *reasons* given by the model for its choices were markedly improved. It was able to justify decisions based on the interrelations between the properties described, including the constraint that they were all contiguous with one another along the road, as was evident from the textual descriptions. The geometry produced after the relational step (Figure 4) also showed improvements:

- The anomalous protrusion was eliminated.
- The displacement issue was partly resolved, as the bearings were fitted to the road.
- GPT-5 correctly identified that parcel 12/21 appeared larger in the cadastral map than the documentary evidence supports — a finding independently confirmed by the formal survey.

The model was effectively correct on the information within its reach, and it also correctly externalised the uncertainty it could not resolve by highlighting it. The residual displacement primarily stemmed from the anchor: 12/21 did not (originally) extend all the way to the south-east intersection as assumed.



Figure 4. Result of Experiment 2 (relational reasoning, single south-east anchor). GPT-5 polygons for parcels 12/38, 12/20 and 12/21 in blue; cadastral map boundaries in red. Residual discrepancies are attributable to an inconsistency between the source documents and the choice of anchor point.

4.3 Experiment 3: Anchor revision — switching to the south-west corner

GPT-5's own relational analysis in Experiment 2 identified a discrepancy between the documented position of the south-east anchor in the *skyltskifte* and its apparent position in the

aerial photograph. Experiment 3 therefore repeated the Experiment 2 setup identically but with the anchor moved to a road corner at the westernmost point in the cluster, where the cadastral records and the aerial photograph agreed well. Only this single anchor was supplied; all remaining geometry was expected to be inferred relationally from the *skyldskifter* and the photograph.

The result (Figure 5) showed further improvement: the cluster geometry aligned more closely with the road and fence lines visible in the aerial photograph, and GPT-5 correctly explained the improvement in terms of the anchor change, noting that the south-east corner had introduced a systematic bias propagating through the entire relational chain. The model also highlighted a remaining discrepancy in the lengths of the boundaries around 12/38 by showing two possible alternatives that would both have been reasonable in the absence of any additional information.

This highlights a practically important property of the puzzle strategy: once a suitable anchor is identified, the model's relational reasoning propagates it accurately across the entire cluster.



Figure 5. Result of Experiment 3 (anchor revised to south-west corner). GPT-5 boundary proposals (blue) for the Gnr. 12 cluster; cadastral map boundaries in red. The cluster geometry aligns more closely with the road and fence lines.

4.4 The model as cadastral inspector

Across all three experiments, GPT-5 repeatedly identified inconsistencies and ambiguities in the source documents that had not been previously noted, including: a documented area for parcel 12/21 inconsistent with the boundary measurements; a polygon that could not be closed, probably because lengths were inaccurately reported (apparently it was common to use step length in place of precise measurement on Røst); and compass bearings unlikely to be precise. This functionality — auditing historical documents and flagging important interconnections, inaccuracies and inconsistencies for expert review — may be as practically useful for the working surveyor as the boundary reconstruction itself.

5. THE PUZZLE STRATEGY: A HOLISTIC APPROACH TO BOUNDARY RECONSTRUCTION

The relational-reasoning approach above can be generalised into what we call the puzzle strategy: treating the parcels as pieces of a jigsaw that must fit together consistently.

5.1 Core steps

1. **Define the analysis area.** Identify the cluster of parcels to reconstruct simultaneously. Each parcel should be bounded primarily by others in the cluster, by known permanent features, or by already known boundaries.
2. **Build the parcel-formation tree.** Extract the subdivision history from the document corpus. The resulting directed graph imposes hard geometric containment constraints and determines the reconstruction order: oldest first.
3. **Collect anchors and known features.** Identify a minimal set of spatial objects that appear in the documents and can be geocoded (road intersections, cadastral benchmarks, coastline segments). A single well-chosen anchor can substantially reduce displacement error for the whole cluster.
4. **Solve the puzzle via structured prompting.** Present all documents, the formation tree and the anchors to the LLM simultaneously. First ask it to analyse spatial relationships and formation history in natural language; then ask it to generate all parcel geometries. *Iterate*: think about inconsistencies with the model, ask it to explain and resolve them, and request a revised output, then repeat this process until the result seems reasonable.

In the final step, we found it useful to display preliminary polygons in a map viewer during the prompting loop, revising a concrete proposal for the geometry of the properties in question in collaboration with the LLM.

5.2 An argumentation theoretic perspective on puzzling

Historical cadastral documents describe boundaries in relational terms. Processing documents one at a time discards this relational information, so that approach seems unsuitable in general and especially so in the Norwegian setting. A deeper reason why the holistic approach should be expected to outperform single-document processing lies in how it allows us to address the problems of incompleteness and inconsistency. When a single document is internally incomplete or inconsistent — as with the non-closing polygon in Experiment 1 — the model has no principled basis for choosing among a potentially large set of geometrical completions. The best we can hope for is that it will choose something reasonable, but as our case study shows, even this cannot be guaranteed. Moreover, even if a reasonable choice is made, it is still likely to be an arbitrary one, unless more information is provided.

The puzzle strategy transforms this situation: incompleteness becomes less likely, since there is more information available, provided we are able and willing to use this information. Meanwhile, there may well arise *more* inconsistencies than before, since we are using much more relational and uncertain information than in a single parcel setting. However, inconsistencies now present themselves as conflicts between evidence items, and the broader evidentiary record together with the applicable legal framework can facilitate a principled resolution. Indeed, when actual property disputes are resolved, high-quality reconstructions are carried out precisely because there are visible inconsistencies that require a resolution. Identifying such inconsistencies and resolving them in advance, before they become conflicts between humans, is a feature of the puzzle strategy, not a bug.

This then also becomes a structured form of a type of reasoning studied in formal argumentation theory (Dung 1995; Bench-Capon and Dunne 2007). In Dung's abstract argumentation framework, arguments attack one another and a preferred extension — the set of mutually consistent, maximally supported claims — is identified. Applied here, each piece of boundary evidence is an argument; conflicting measurements are attacking arguments; and the legally and evidentially best-supported resolution is the winning extension.

Walton et al. (2008) further show how argumentation schemes — structured patterns of defeasible inference — can guide conflict resolution in domains governed by legal texts, which also maps naturally onto our setting. Such a structured, legally grounded approach to inconsistency resolution is facilitated when we adopt the puzzle strategy instead of single-document workflows, and LLMs, trained on legal and technical texts, seem well suited to execute it. In future work, we would like to explore semi-automated argumentation further, as a possible pathway to a systematic approach to inconsistency resolution in this setting.

The iterative prompting structure — reasoning before geometry, then explicit flagging of inconsistencies — also mirrors the mixed-initiative interaction model (Horvitz 1999): the surveyor sets goals at a high level of abstraction, the AI handles detailed cross-referencing, and the system cycles between them until a satisfactory result is reached. This is also analogous to the ReAct (Reasoning and Acting) agent paradigm of Yao et al. (2022), where the model alternates between explicit reasoning steps and concrete actions. Indeed, it might soon be possible to automate parts of the prompting currently done by a human, if we succeed in developing agent-based systems where LLMs engage in critical dialogue with one another about property boundaries to more accurately recreate them.

6. DISCUSSION AND CONCLUSION

The iterative human-in-the-loop workflow demonstrated in this paper indicates the future potential of multi-modal LLMs for cadastral work in Norway. It is not yet possible, neither in practice nor by law, to fully delegate cadastral surveying to AI systems. However, the AI can legitimately produce preliminary boundary proposals that the surveyor verifies, analyse and

summarise documentary evidence, and flag inconsistencies that direct field resources to the boundary segments most in need of physical measurement — none of which requires any change to the current legal framework. This strategy is also not Norway-specific: it depends only on a document corpus describing spatial relationships among parcels and a minimal set of anchor coordinates, both of which are typically present in challenging cases of boundary reconstruction.

This also highlights an underappreciated role for LLMs in cadastral work: as document inspectors that can help the surveyor analyse a large volume of historical records that may have escaped in-depth human attention for decades. Indeed, this may be the most immediately deployable path to practical value, facilitating a workflow where large amounts of data may be accessed through an LLM that can assist with the analysis and point to key data points.

Key directions for future work include:

- Automated parcel-formation tree construction, necessary for scaling to large document corpora.
- A lightweight constraint solver to formalise and validate the spatial reasoning used by the LLM to recreate the boundaries based on the documentation.
- A fully agentic architecture where the LLM iteratively calls GIS tools to compute and verify its own output, removing the human prompter from the inner loop.
- Large-scale field validation with formal survey results as ground truth.

A working research prototype implementing the puzzle strategy as a surveyor copilot — combining an interactive map, a chat interface, and an LLM agent with access to geocoding and GIS tools — is currently under development at HVL.

REFERENCES

- Bench-Capon, T.J.M., Dunne, P.E., 2007. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10–15), 619–641.
- Dung, P.M., 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2), 321–357.
- Horvitz, E., 1999. Principles of Mixed-Initiative User Interfaces. Proc. CHI 1999. <https://doi.org/10.1145/302979.303030>.
- Larsson, G., 1991. *Land Registration and Cadastral Systems: Tools for land information and management*, Addison Wesley Longman Limited.
- Leiknes, A., Mjøs, L.B., 2021, Informasjon om eksisterende eigedom, In Røsnes, A. (Ed), *Arealadministrasjon* (2. ed), 129-157. Universitetsforlaget.

Liu, N.F., Lin, K. et al 2024. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.

Walton, D., Reed, C., Macagno, F., 2008. *Argumentation Schemes*. Cambridge University Press.

Yao, S., Zhao, J., et al., 2022. ReAct: Synergizing Reasoning and Acting in Language Models. arXiv:2210.03629.

Recreating property boundaries from old maps and cadastral records using AI: The case of Røst, Norway (14110)
Sjur Dyrkolbotn, Arve Leiknes and Leiv Bjarte Mjøs (Norway)

FIG Congress 2026
The Future We Want - The SDGs and Beyond
Cape Town, South Africa, 24–29 May 2026