Cooperation Atlas (CoopAt)

Austrian Science Fund (FWF) Project Proposal

Applicant:	Univ.-Prof. Dr. Dr. Ulrich Frey
Institution:	University of Graz
Duration:	36 months
Date:	August 28, 2026

Contents

1 Main Objective

Research on common-pool resources and social-ecological systems (SES) has generated influential theories, extensive case-study evidence and several comparative databases, yet this knowledge remains fragmented in three mutually reinforcing ways. Empirical evidence is dispersed across studies and databases whose variables, case boundaries and outcome definitions cannot be compared directly [1]. Different SES frameworks make different aspects of the same case visible [2]. Comparison itself forces a trade-off between the context a case study preserves and the standardisation that large-scale analysis requires [3, 4]. A fourth obstacle is independent of these three: with the data unchanged, different model classes lead to different conclusions about which factors matter [5].

CoopAt asks which configurations of social, institutional and ecological factors explain success in forests, fisheries, irrigation systems and other user-governed commons, and under which conditions. It is aimed at what the field disputes rather than at what it takes as settled. Congruence between rules and local conditions and proportionality between investment and extraction appear to hold across resource types, whereas the weight of monitoring and graduated sanctions shifts with the resource itself: both matter most where it is static enough to be observed, while clearly defined social boundaries matter more where it is mobile [6]. These questions will be examined across resource domains rather than within one, under multiple theoretical frameworks and on a far larger body of case evidence than any existing comparative database holds, a combination none of them was designed for. The cost of expert hand-coding has bounded both how many cases could be compared and how much context survived the coding, and maintaining the resulting databases has been a persistent obstacle in its own right [7]. Case-study documents will therefore be processed into structured, source-linked records by LLM-assisted extraction under expert oversight [8], and every coded value will retain the passage it came from, the coding decision behind it and its uncertainty. A result counts where it survives changes in framework, contextual information and analytical method, including the shift from estimating average factor effects to identifying sufficient combinations of conditions; where it does not, the divergence is reported as a finding rather than settled by the choice of method. Success is analysed in two dimensions, ecological condition and socio-economic benefits, together with the interactions between them [1, 9].

The primary contribution is an account of how far explanations of success depend on the theoretical perspective through which the evidence is read, a question the field has not been in a position to ask and one that is informative in either direction. Where frameworks agree, a finding is stronger than any one of them could establish alone, and where they disagree the project will show whether they encode genuinely different observations or only different vocabulary. This bears on the long-standing division between qualitative case-study research and quantitative large-N comparison, where the methods for combining case-oriented with variable-oriented analysis are themselves underdeveloped [10]. CoopAt addresses both the infrastructure limit and the methodological one, and the corpus, the extraction workflow and the data model are released so that the comparison can be repeated, contested and extended beyond this project.

2 Introduction

Forests, fisheries, irrigation systems and other shared natural resources support livelihoods and collective well-being, but their long-term governance poses a recurrent problem [9, 11]. In common-pool resources, one person's use can reduce what remains for others, while excluding potential users is difficult or costly. Individually reasonable decisions can therefore undermine both the resource and the cooperation needed to maintain it [12]. Resource degradation is not inevitable, however: communities can develop durable institutions for governing shared resources [13]. The difficult question is not whether collective governance is possible, but which arrangements work for which resources, users and circumstances.

The commons and SES research tradition approaches this question diagnostically rather than prescriptively. Rather than prescribing one best institutional design, it decomposes a case into components whose interactions can be examined. One of the most influential decompositions in commons research distinguishes resource systems, resource units, actors, governance arrangements, interactions, outcomes and the wider social, economic and ecological setting [11], which makes it possible to ask why similar institutions produce different results, or why different arrangements succeed under apparently similar conditions. It is not, however, the only decomposition available. A comparison of ten established SES frameworks, spanning ecosystem services, vulnerability, sustainable livelihoods, material and energy flows, adaptive water governance and Ostrom's own, found three criteria sufficient to distinguish them: whether the social and ecological systems are conceptualised as linked in one direction or reciprocally, whether the ecological system is viewed in anthropocentric or ecocentric terms, and whether the framework is built for analysis or for intervention [2]. Each makes different aspects of the same case visible. The choice between them is not neutral bookkeeping: it determines which factors are recorded and, in consequence, which explanations of success can be formulated at all.

The outcome to be explained is itself contested. A resource can remain ecologically productive while benefits are distributed unequally, or governance can become more inclusive while resource pressure increases. Success may refer to ecological condition, sustainable resource use, socio-economic benefits, governance quality or equity, and assessments of it differ among actors, organisational levels and spatial or temporal scales [9]. Each additional dimension has to be coded consistently across every case, and the more a case record is required to carry, the less comparable data remain [3]. CoopAt therefore analyses two dimensions, ecological condition and socio-economic benefits, and treats the interactions between them as part of what is to be explained rather than as a complication, because gains in one can coincide with losses in the other and because aggregating outcomes into a single score obscures how a given factor bears on each of them [9]. The question is then not what makes an SES successful, but which conditions matter for which dimension of success, and where the two diverge.

Against this background it is worth being precise about what the field has settled and what it has not. Congruence between rules and local conditions, and proportionality between what users invest and what they extract, recur across resource types. The weight of monitoring and graduated sanctions, by contrast, shifts with the resource: both matter most where it is static enough to be observed, while clearly defined social boundaries matter more where it is mobile [6]. Whether group size and user heterogeneity help or hinder collective governance has likewise resisted a general answer [14]. Findings of this kind shift explanation away from isolated variables: monitoring, participation or trust may matter, but their effects depend on how they combine with each other and with the characteristics of the resource and its users [15]. What has to be explained is therefore a configuration of conditions, and comparative research must seek general patterns without treating context as noise.

SES and commons research has accumulated an extensive and richly detailed body of case-based evidence. The challenge is not a lack of knowledge about individual systems: case studies preserve mechanisms and local context, but differences in the variables examined and in how concepts are defined and operationalised make systematic comparison across contexts difficult [7, 15]. Comparative databases improve standardisation, but their different research histories, sectoral purposes and coding architectures limit direct integration. Definitions and measures of success add another source of inconsistency [9]. Finally, analytical method affects the conclusions drawn from a dataset. Using the same irrigation data, Frey [5] found both agreement and substantial differences in estimates of factor importance across statistical and machine-learning models, despite broadly similar predictive performance. A finding can therefore appear convincing within one analysis while remaining ambiguous when reasonable alternative models produce different substantive conclusions. Schlüter et al. [4] identify these difficulties as a gap in the field's methods for synthesis and theory building. Generalising across accumulated case knowledge is held back by the absence of approaches able to handle different types of data and the diversity of concepts and methods that produced them, and the comparative databases built in response buy standardisation at the cost of fitting specific contexts. They also note that every method carries its own limitations and biases, which is why they call for multi-method approaches and triangulation rather than reliance on any single one.

Together these limitations prevent the field from accumulating reliable knowledge about what works under which conditions. Researchers can identify many plausible influences on SES outcomes, but it remains difficult to determine which relationships generalise, which apply only to particular resource domains or regions, and which are artefacts of a chosen analytical method. This also limits the guidance available to communities, resource managers and policy makers, who need to understand when an institutional arrangement is likely to fit a particular setting rather than receive another universal prescription. A result counts, then, where it survives changes in the framework through which cases are represented, in the contextual information retained and in the analytical method applied; where it does not survive, the divergence is informative in its own right and should be reported rather than settled by the choice of method [5, 7, 15].

Meeting that standard makes demands on the evidence that a table of coded values cannot satisfy. Analytical approaches differ in what they need: set-theoretic comparison needs values calibrated against close familiarity with each case, statistical and machine-learning models need dense variable matrices, and relational or simulation-based approaches need other representations again. A database of hand-coded values serves one of these, because the reasoning that produced each value is no longer available; a corpus in which every coded value retains the passage it came from can be read again for the others. Building such a corpus has until now been limited by the cost of expert hand-coding, which bounded both how many cases could be compared and how much context survived the coding, and by the effort of maintaining the resulting databases [7].

CoopAt brings these requirements together. The applicant's 24-factor framework provides the coding structure for harmonising existing databases and for guiding extraction [15], while the source passage retained behind each coded value allows the same cases to be represented under other frameworks and analysed by methods with different requirements. Case-study documents are processed into structured, source-linked records by LLM-assisted extraction under expert oversight [8], with the coding decision and its uncertainty recorded alongside each value. CoopAt analyses what case-study documents record rather than observing systems in the field, so processes that unfold over time enter the analysis only as far as the underlying sources describe them. The framework, the corpus and the extraction workflow are the means; the intended result is an account of which conditions matter, for which dimension of success and under which circumstances, and of how far that account depends on the perspective through which the evidence is read.

3 State of the Art

CoopAt draws on four bodies of research: diagnostic SES frameworks, comparative commons databases, LLM-assisted information extraction and multi-method modelling. Together, these bodies of research define both the foundations of the project and the unresolved problem it addresses.

3.1 SES and Commons Research

A central development in commons research was the move from universal prescriptions towards a diagnostic approach to institutional diversity. Ostrom's Institutional Analysis and Development (IAD) framework established this mode of analysis, and the later SES framework developed a shared vocabulary that distinguishes resource systems, resource units, actors, governance arrangements, interactions and outcomes [11, 12]. These frameworks make heterogeneous cases comparable while retaining the premise that the effects of governance arrangements depend on their social and ecological setting.

Subsequent comparative studies found recurring support for institutional features such as monitoring, graduated sanctions, accessible conflict-resolution mechanisms and rules adapted to local conditions, while also showing that their effects depend on the social-ecological context [16, 17]. The applicant's subsequent synthesis extended this work into a 24-factor framework spanning resource systems, resource units, actors, governance systems and external influences. It locates comparison at a practical middle level between concepts that are too abstract to measure and variables that are too case-specific to compare, but it was organised primarily around sustainability and identifies economic indicators, wider contextual conditions and additional outcome dimensions as areas for further development [15]. The unresolved task is therefore to refine and operationalise the framework so that its factors can be tested consistently across resource domains and multiple dimensions of success.

3.2 Existing Databases and Framework Limitations

Comparative database projects have already shown that detailed commons evidence can support larger cross-case analyses. The Common-Pool Resource (CPR), Nepal Irrigation Institutions and Systems (NIIS) and International Forestry Resources and Institutions (IFRI) initiatives, followed by the Social-Ecological Systems Meta-Analysis Database (SESMAD), translated evidence from different sectors and research traditions into structured case records [18]. These projects created valuable datasets and reusable coding protocols, but they were designed around different questions, case boundaries and units of analysis. As a result, related concepts are often represented in forms that cannot be combined directly, and the datasets make different tradeoffs between contextual detail and comparability [7].

These differences are scientifically consequential, not merely technical. Pooling variables that use different definitions or levels of analysis would increase the number of cases without making them genuinely comparable. Conversely, imposing a uniform coding structure risks losing the case-specific information needed to interpret those values, reflecting a longstanding tension between increasing the number of observations and preserving qualitative depth [3]. CoopAt therefore needs a harmonisation strategy that translates existing variables into the refined framework while retaining the reasoning and uncertainty behind each coding decision. Only then can the larger dataset provide a defensible basis for testing whether different models converge on which factors matter under which conditions.

3.3 LLM-Assisted Information Extraction

In SES and commons research, information extraction is not a simple term-matching task because variables are coded from narrative case studies, often with ambiguous evidence, missing information and coder judgement [19]. A central bottleneck for comparative SES databases is their long-term maintenance, because technical upkeep, new case entry and detailed coding often extend beyond short project funding cycles [7]. CoopAt addresses this bottleneck by testing whether an LLM-assisted extraction pipeline can reduce the effort needed to extend SES databases while preserving traceable evidence and expert review.

Recent information-extraction research shows that LLMs can convert unstructured text into structured knowledge such as entities, relations and events [20]. In scientific domains, schema-constrained LLM extraction has also been shown to produce normalised records for database construction [8]. Recent multimodal LLMs extend this shift beyond text-only extraction, because they can work with image and text inputs, while scientific figures are increasingly treated as a distinct multimodal comprehension problem rather than as plain text alone [21–23]. For CoopAt, however, the relevant benchmark is not extraction in general but framework-guided SES coding.

Work on LLM-assisted qualitative coding and social-science annotation shows that performance varies with codebook design, prompt formulation, model choice and task complexity, and that outputs must be evaluated against human-coded samples [24–26]. For this reason, LLM-based extraction must be treated as a validated coding workflow with documented prompts, model-version control, source passages, confidence information and checks for bias, instability and systematic failure [26, 27]. The open challenge for CoopAt is to adapt these advances to SES case evidence by linking each extracted value to the 24-factor codebook, the source passage and human review where needed.

3.4 Predictive Modelling with SES Data

Quantitative applications of the SES framework have expanded, but recent reviews show that methods remain heterogeneous and that many studies still struggle to capture interactions among SES indicators [28]. Earlier work using the applicant's 24-factor system showed that neural networks can model ecological success from coded CPR and SES data, while also underlining that such models depend on carefully constructed variables and outcome measures [29, 30]. Later multi-method modelling with NIIS data showed that estimates of factor importance can differ substantially across GLM, random-forest, gradient-boosting and neural-network approaches, even with the same underlying dataset [5].

Since these studies, quantitative SES research and machine-learning practice have developed further, making it timely to revisit the 24-factor modelling approach with a larger harmonised dataset and a model-comparison strategy that treats robustness as an empirical question. For CoopAt, modelling is therefore not a replacement for qualitative interpretation, but a way to test whether patterns identified across many cases are robust to model choice, missing data and domain differences. The project will compare interpretable statistical models with more flexible machine-learning methods, so that non-linear interactions can be explored without treating prediction accuracy as the only criterion.

4 Project Aims

CoopAt pursues four research questions (RQs):

RQ1: **Success patterns:** Which factors and factor configurations are robustly associated with different dimensions of SES success across model classes, outcome definitions, resource domains, regions and data sources, and where do synergies, trade-offs or context-specific limits appear?

RQ2: **Evidence base:** How should the applicant's 24-factor SES framework, outcome coding, data model and documented API be refined to integrate legacy databases and case-study literature while preserving comparability, context and reuse by other researchers?

RQ3: **Maintainable extraction:** How reliably can an LLM-assisted extraction pipeline recover SES factors, outcome definitions and evaluative criteria from case-study literature, and where does it require human review?

RQ4: **Primary-data enrichment:** How can framework-guided primary data collection validate, complete and contextualise extracted and harmonised SES evidence in contemporary resource-governance cases?

The central hypothesis is that a larger harmonised dataset, a refined multi-factor framework and more advanced machine-learning models will improve inter-model agreement under systematic robustness checks across model families. The substantive expectation tested within this comparison is that different dimensions of SES success are associated with identifiable factor configurations rather than with isolated factors alone, and that these relationships vary across resource domains, regions and data sources. Refined factor definitions, outcome-coding rules and a data model that records source and processing history are expected to make heterogeneous SES evidence more comparable without removing the context needed for interpretation and reuse. A documented LLM-assisted extraction pipeline is expected to reduce the practical barriers to maintaining and extending SES databases, while reaching expert-comparable reliability for clearly defined factors, outcome definitions and evaluative criteria and requiring human review for ambiguous or context-dependent variables. Framework-guided primary data collection is expected to fill evidence gaps, validate ambiguous factors and contextualise model results with contemporary resource-governance cases.

The research questions and working hypotheses are translated into five operational aims in Table 1. The table links each aim to the data sources, analytical steps, outputs and success criteria that will be used to evaluate progress. The detailed timing and dependency logic are presented in Section 8.

Aim	**Method and data**	**Analysis**	**Output and success criterion**
Refine the 24-factor SES framework	Test factor definitions, criteria, indicators and outcome dimensions against legacy variables, representative case texts and expert feedback.	Record ambiguous cases and revise lower-level criteria or indicators only when comparability improves.	Codebook version 1 by month 6 and revised codebook by month 12, with success defined as mapping and outcome-coding rules usable for databases and extraction prompts.
Harmonise existing SES databases	Map accessible sources such as CPR, NIIS, IFRI, SESMAD and related datasets into the framework.	Track original definitions, outcome concepts, scales, recoding decisions, mapping confidence and unmappable variables.	Publicly inspectable mapping documentation by month 12 and harmonised dataset version 1 by month 18, with success defined as traceable mapping rather than forced completeness.

Aim	Method and data	Analysis	Output and success criterion
Develop and validate LLM extraction	Process case-study documents with prompts based on the codebook, source evidence, confidence scores and uncertainty notes.	Compare LLM-assisted outputs with human-only extractions in random validation samples, rated by experts without method labels.	Validated extraction workflow by month 18 and literature dataset by month 24, with success defined as factor-level and outcome-level evidence on where LLM assistance is reliable.
Integrate and model multidimensional SES success	Combine harmonised, extracted and enriched data in analysis-ready data frames with documented outcome construction and recorded source and processing history.	Compare interpretable statistical models, including generalized linear and mixed-effects models, with random forests, gradient boosting and neural networks, using model-family, domain and region robustness checks.	Integrated CoopAt dataset by month 24 and scientific findings and robustness report by month 30, with success defined as distinguishing robust patterns from model-dependent findings, not headline prediction alone.
Validate and disseminate CoopAt	Use AIIS partner input, expert review, documentation and open dissemination.	Check whether the framework captures relevant variation in contemporary irrigation and resource-governance cases.	Visualisation, API documentation and replication package by month 36, with success defined as reusable findings and empirical infrastructure for future SES research.

Table 1. Operationalisation of project aims, outputs and success criteria.

4.1 Risk Assessment and Fallbacks

Table 2 summarises the main implementation risks and the fallback strategies that keep the work plan viable. The risks are formulated as project-management checks rather than expected failures.

Risk	Mitigation
The 24-factor framework proves too coarse or too ambiguous for reliable classification.	Use the first project phase to refine factor definitions, criteria, indicators and mapping rules against real legacy variables and case texts. Split or merge lower-level criteria or indicators only where this improves expert agreement and model usability.
Legacy databases cannot be fully accessed or mapped.	Prioritise accessible sources first, document unmappable variables explicitly, and compensate with LLM-extracted literature cases and AIIS validation material.
LLM extraction accuracy is insufficient for some factors.	Retain human-in-the-loop review for low-confidence variables, use simpler factor-specific prompts, and restrict automated extraction to variables that meet validation thresholds.
Primary-data enrichment is delayed by partner availability.	Treat primary data as validation rather than a critical dependency and proceed with harmonised legacy data and literature extraction while using partner input for expert review.
Models produce unstable or overgeneralised claims about SES success.	Use cross-validation by domain and region, compare interpretable statistical models with more flexible machine-learning models, report model-dependent findings explicitly, and distinguish robust patterns from exploratory or domain-specific results.

Table 2. Risk assessment and fallback strategies.

5 Originality and Innovation

CoopAt is original because it combines three elements that are usually treated separately: a balanced SES framework, large-scale extraction from scientific case-study documents, and comparative modelling of multidimensional SES success across resource domains. Existing SES research has strong conceptual foundations, but the field still lacks a data infrastructure that is sufficiently comparable, fine-grained and scalable for robust analysis of factors, configurations and tradeoffs. The project addresses this gap by turning the framework itself into a classification and integration instrument.

The first innovation is conceptual. The 24-factor framework is designed to cover resource-system, resource-unit, actor, governance-system and external-environment factors at a comparable level of abstraction. This avoids a common weakness of SES databases: ecological or biophysical variables may be measured in detail while institutional, actor-related or external-context variables remain broad, or vice versa. By refining the framework against both existing databases and new extraction tasks, the project creates a practical bridge between conceptual SES theory and empirical measurement.

The second innovation is methodological. CoopAt does not develop new multimodal AI models. Instead, it uses recent advances in multimodal LLMs to read SES case studies more completely: not only the prose, but also tables, captions, figures, diagrams and document layout [21, 22]. This matters because important information about actors, resources, governance arrangements, causal assumptions and outcomes is often communicated through several parts of a paper at once. Scientific figures are also increasingly treated as a distinct multimodal understanding problem rather than as plain text alone [23]. The innovation is not the model architecture or the use of an LLM as such, but the controlled workflow around these tools: codebook-based prompts, integrated reasoning over textual and visual evidence, source-linked outputs, confidence values, manual review of uncertain cases and validation against already coded material. This makes automated extraction accountable to the standards of empirical social science.

The third innovation is integrative. CoopAt deliberately avoids both overly rigid database structures that strip away case context and overly flexible qualitative repositories that prevent comparison. The project uses data frames for numerical and categorical variables while retaining text comments and source passages. This gives the modelling workflow access to comparable data while allowing researchers to inspect where coded values came from, how they were transformed and where later recoding may be needed. The linked records of source and processing history are also important for AI-assisted research workflows. Rather than exposing only context-free variables, CoopAt will allow human researchers and AI tools to follow coded values back to source passages, coding notes and uncertainty information. This does not replace qualitative interpretation, but it enables use-case-specific checking, contextualisation and recoding of structured evidence.

The fourth innovation is analytical. By combining GLMs with random forests, gradient boosting and neural networks, the project can compare interpretable baseline models with more flexible predictive models. This will show not only which factors are associated with specific success dimensions, but also where interactions among ecological, institutional and social variables matter and where tradeoffs limit simple success claims.

6 Methods

The project follows a staged mixed-methods design. It begins with framework refinement, proceeds through database harmonisation and LLM-assisted extraction, adds targeted primary-data enrichment for validation, and culminates in integrated modelling and dissemination. All stages serve the substantive analysis of multidimensional SES success. The common methodological principle across all stages is traceability: every structured value and outcome definition must be connected to its source, extraction decision and confidence level. The workflow is dependency-aware: large-scale extraction starts only after the codebook and initial prompts have been tested, and modelling starts only after a stable integrated dataset with documented source and processing history is available.

Framework Refinement

The starting point is the applicant's existing 24-factor system that synthesises SES and commons research [15]. Following this synthesis, the 24 factors are organised as resource system, resource units, actors, governance system and external environment. For example, CoopAt will treat resource size, resource boundaries and initial ecological condition as resource-system factors; manageability and regeneration as resource-unit factors;

number of actors, group composition, social capital and resource dependency as actor factors; and participation, rule characteristics, monitoring, compliance and conflict management as governance-system factors. External-environment factors such as exclusion of third parties, relations with external actors and capacity to adapt to change will be retained where case material provides sufficient evidence. Ecological success is treated separately as an outcome dimension, not as one of the 24 explanatory factors. In the first phase, the framework will be tested against real variables from legacy databases and against representative case-study texts. The team will produce a codebook that keeps the 24 factors as the main organising structure and refines the lower-level criteria, indicators and mapping rules needed for consistent coding. Particular attention will be given to distinctions between variables, criteria and indicators, which must be precise enough for consistent classification but not so narrow that they undermine comparability across resource domains. Each codebook entry will include a definition, inclusion and exclusion rules, examples of clear and borderline cases, expected data type, admissible values where applicable, recoding guidance and guidance on how to record uncertainty.

Legacy Database Harmonisation

Existing datasets such as NIIS, CPR, IFRI and later SES database initiatives will be mapped into the 24-factor framework where access and documentation permit. For each source variable, the team will record its original name, definition, scale, admissible values, missingness pattern and mapping decision. Where possible, source variables will be mapped to indicators or criteria first and then linked to the corresponding factor-level variable. Variables that cannot be mapped cleanly will not be forced into the framework. They will be documented as partial, ambiguous or out-of-scope mappings. The same principle applies to outcome variables: ecological condition, sustainable appropriation, socio-economic effects, governance-process outcomes and related success definitions will be constructed with source-specific rules before they are compared across datasets. This produces both an integrated analysis dataset and a transparent mapping document that can be inspected and revised by other researchers. The mapping document will therefore function as both a data product and a quality-control instrument. Reviewers of the dataset will be able to see which original concepts were preserved directly, which were transformed, and which were excluded.

LLM-Assisted Literature Extraction

The LLM extraction workflow will be developed iteratively. Case-study documents will be prepared as document packages that preserve prose, captions, tables, page references and, where relevant, figure or diagram images. The pipeline will use available multimodal LLMs rather than train new foundation models. Documents will be segmented into manageable evidence units while keeping links to their original location, and prompts based on the codebook will ask the model to reason over textual and visual evidence together where this is needed. For each factor, criterion, indicator, outcome definition or evaluative criterion that can be extracted reliably, the output will include the proposed value, evidence passage, confidence rating and any uncertainty note. Already structured cases will be used to calibrate prompts and assess extraction quality. The workflow will prioritise precision over volume in the first project year and scale up only after the validation results show which factors can be extracted reliably.

Quality assurance has three layers. First, automated checks will flag missing evidence, contradictory values and implausible combinations. Second, an expert-reviewed validation sample will be used to estimate accuracy by factor rather than only at aggregate level. Third, low-confidence and high-impact cases will be manually reviewed. After larger extraction batches, a random sample of cases and variables will be selected for post-hoc validation. For this sample, human-only extractions will be created independently of the LLM-assisted results. Experts will then rate the two extraction results against each other without being told which result is LLM-assisted and which is human-only. If the LLM-assisted results do not reach comparable quality for a factor, that factor will remain subject to human extraction or review until the procedure is improved. The final dataset will therefore distinguish expert-generated, LLM-extracted and human-reviewed values instead of treating all entries as equally certain. It will also record whether an extracted value is supported by prose, tables, captions, figures, diagrams or combinations of these evidence types, without treating text and visual extraction as separate workflows.

Validation Measures

The project will measure extraction and harmonisation quality explicitly. For database mapping, the main indicators are the number of variables and outcome definitions mapped, the share mapped unambiguously, the share requiring partial or qualified mapping, the number excluded with documented reasons and the completeness of recoding and outcome-construction documentation. For LLM-assisted extraction, the main indicators are factor-level agreement with human-only extraction, agreement at the more detailed criterion, indicator or outcome-definition level where applicable, expert preference in blind comparison samples, completeness of source and processing history, contradiction rates, and the proportion of cases routed to manual review. For modelling, the main indicators are out-of-sample performance, robustness across domains and regions, sensitivity to missing data, sensitivity to plausible extraction-error rates and stability of dimension-specific findings.

Primary-Data Enrichment

Primary-data collection is included as a focused validation and enrichment component. The project will build on original Ostrom questionnaire material, later questionnaire developments and the applicant's connections to the Asian Irrigation Institutions and Systems initiative (AIIS). The aim is not to create a separate fieldwork-heavy project, but to test whether the 24-factor framework and its mapping rules capture relevant variation in contemporary irrigation systems and to close selected data gaps where partner knowledge is available. Potential cooperation countries include India, the Philippines, Indonesia, Cambodia and Bangladesh, with Nepal as a possible validation context depending on partner access. Any questionnaire or interview activity involving identifiable persons will be conducted only after the required institutional review, as described in Section 10.

Data Integration

The integrated dataset will use data frames with one row per case and columns for framework variables, outcome dimensions, metadata, and records of source and processing history. Numerical and categorical entries will be accompanied by text comments, source passages and transformation notes that document the immediate evidential context of each structured value. This does not preserve the full qualitative case context, but it keeps the information needed to check, interpret and revise each structured value. This design is intentionally simpler than a fully relational database. It supports statistical modelling and reproducible analysis while avoiding a database structure that would require premature decisions about complex case relationships. The dataset will retain versioned metadata so that later users can identify whether a value came from a legacy database, an LLM-assisted extraction, a human-only extraction, a reviewed correction, or partner-provided enrichment.

Modelling Strategy

The modelling stage will estimate how SES factors relate to different dimensions of sustainability and governance outcomes. Generalized linear models and mixed-effects models will provide the interpretable statistical modelling layer. Mixed-effects models will be used where cases are grouped by resource domains, regions, source databases or studies. This allows the project to estimate factor-outcome relationships while taking such group differences into account. Random forests and gradient boosting will be used to capture non-linear effects and rank variable importance. Neural networks will be explored where the available sample size and data quality justify more flexible modelling. This design follows from prior multi-method work showing that method choice can change which natural-resource-management factors appear most important, so agreement and disagreement among model classes will be treated as a substantive finding rather than as a purely technical detail [5]. Model inputs will distinguish original variables, harmonised indicators and factor-level aggregates so that modelling choices remain auditable. Missing values will be handled through transparent imputation and sensitivity checks, and all results will be reported with uncertainty estimates. Where the data permit, the project will compare models across outcome dimensions to identify synergies, tradeoffs and domain-specific limits. Composite success profiles will be used only after the individual dimensions have been inspected and documented.

Validation will combine cross-validation, domain holdouts and region-based robustness checks. Models trained on one resource type, such as irrigation systems, will be tested against other domains, such as forests or fisheries, where the data permit. A finding will be treated as robust only when its substantive direction or importance remains stable across outcome construction choices, model families, missing-data treatments and domain or region holdouts. The project will emphasise explanatory robustness over headline prediction accuracy: a model that performs well only in one domain or for one success dimension will be treated as a domain-specific or dimension-specific result, not as a general SES model.

7 Cooperation Arrangements

CoopAt is designed to be feasible as a University of Graz project led by the applicant and the funded project team. The core tasks–framework refinement, database harmonisation, LLM-assisted information extraction, data integration and modelling–can be completed within the project group. External cooperation is therefore used to improve empirical coverage, validation quality and dissemination rather than to create dependencies that would endanger project delivery.

The most important cooperation context is the applicant's involvement in AIIS. AIIS connects researchers and field partners working on irrigation systems in Asia using Ostrom-inspired concepts and questionnaire material. For CoopAt, this network provides access to expert feedback on the 24-factor framework, opportunities for targeted primary-data enrichment, and validation cases in countries such as India, the Philippines, Indonesia, Cambodia and Bangladesh. Any primary-data activity will be scoped as enrichment and validation of the framework, not as the main empirical basis of the project.

The project also draws on the applicant's existing links to commons scholarship, including the International Association for the Study of the Commons (IASC) and the Ostrom Workshop community. These networks are relevant for expert review of the codebook, identification of legacy data sources and dissemination of the resulting CoopAt infrastructure. Where access to existing databases depends on individual custodians or author contacts, the project will first use publicly available documentation and then seek clarification or missing information through direct scholarly contact.

Within the University of Graz, the project is embedded in Environmental Systems Sciences and the *Complexity of Life* research profile area. This setting provides interdisciplinary expertise in sustainability, systems science, modelling and data-driven environmental research. The university's research service, library, information technology (IT) and high-performance computing support will contribute to the practical implementation of the extraction and modelling workflow.

8 Work Plan and Timeline

The project is planned for 36 months. The work packages are sequenced so that framework refinement and outcome coding begin immediately, data harmonisation and LLM extraction are validated before they scale, and modelling starts once a sufficiently stable integrated dataset is available. Milestone (M) and deliverable (D) labels are used consistently in the work-package descriptions. Table 3 summarises the main phases and shows how each phase depends on earlier outputs.

Phase	Months	Dependency logic
Framework and project setup	1–12	Codebook, mapping rules and documentation standards are established before full database mapping and extraction scale-up.
Database harmonisation	3–18	Legacy mapping documents provide training, validation and baseline data for the extraction workflow.
LLM extraction	6–24	Prompt development begins early, but large-scale extraction follows validation by factor and, where possible, by indicator.
AIIS enrichment and expert validation	12–30	Partner input validates selected variables and does not block core dataset construction.
Integration, scientific modelling and dissemination	18–36	Modelling starts after an integrated dataset exists and is updated through final validation.

Table 3. Compact dependency-aware timeline.

8.1 Work Package 1 - Project Management and Framework Refinement

Start: Month 1 **End:** Month 36

Goals

- Establish project routines, documentation standards and ethical oversight.
- Refine the 24-factor SES framework and produce a project codebook with mapping and outcome-coding rules.
- Coordinate integration across data, extraction and modelling tasks.

Content

The first work package combines project management with the conceptual groundwork needed by all later tasks. The team will set up version-controlled documentation, classification guidelines, meeting routines and a risk register. In parallel, the 24-factor framework will be tested against variables from legacy databases and representative case-study texts. The output is a stable codebook with factor definitions, outcome dimensions, criteria, indicators, classification examples and rules for ambiguous cases.

Milestones and Deliverables

- M 1.1 Project infrastructure and documentation workflow established (month 2).
- D 1.1 24-factor SES codebook with initial mapping and outcome-coding rules completed (month 6).

- D 1.2 Revised codebook after database and extraction feedback (month 12).
- M 1.2 Risk register and data ethics checklist reviewed at each six-month project review.

8.2 Work Package 2 - Existing Data Sources and Harmonisation

Start: Month 3 **End:** Month 18

Goals

- Identify, document and prioritise existing SES databases and case collections.
- Map accessible variables from legacy databases into the 24-factor framework.
- Produce transparent mapping documentation between original variables, indicators and CoopAt variables.

Content

This work package integrates existing data sources such as NIIS, CPR, IFRI and later SES database initiatives where access and documentation permit. The team will document original variable definitions, scales, missingness patterns and recoding decisions before mapping them into the framework. Cases or variables that cannot be mapped reliably will be retained as documented limitations rather than forced into the analysis dataset. Outcome variables will receive the same treatment, so that ecological, socio-economic, governance and related success dimensions are constructed transparently before comparison. The harmonised data form the first empirical layer of CoopAt and provide training and validation material for LLM-assisted extraction.

Milestones and Deliverables

- D 2.1 Inventory of accessible databases and variables (month 6).
- D 2.2 Database-to-framework mapping documentation (month 12).
- M 2.1 Harmonised legacy dataset version 1 available for modelling and validation (month 18).
- M 2.2 Unmappable variables documented with reasons rather than omitted silently (month 18).

8.3 Work Package 3 - LLM-Assisted Information Extraction

Start: Month 6 **End:** Month 24

Goals

- Develop extraction prompts based on the factor, outcome-dimension, criterion and indicator codebook.
- Validate extraction quality against already structured and expert-reviewed cases.
- Implement post-hoc validation with random samples and expert rating of anonymised extraction results.
- Scale literature extraction only for factors that meet reliability thresholds.

Content

The project will process selected SES case-study literature through an LLM-assisted workflow. The workflow will use available multimodal LLMs to reason over prose, tables, captions, figures and diagrams where these jointly support a coding decision. Each extraction will return a structured value, supporting source evidence and confidence rating for SES factors, outcome definitions and evaluative criteria. For random validation samples, the team will create human-only extractions independently of the LLM-assisted results. Experts will then rate anonymised extraction results against each other, measure agreement and accuracy by factor and by more detailed indicators where applicable, and route low-confidence or inconsistent outputs to manual review. The work package will prioritise validated, traceable extraction over maximum throughput.

Milestones and Deliverables

- D 3.1 Extraction prompt set and processing workflow aligned with the codebook (month 12).
- M 3.1 Validation sample extracted and accuracy assessed by factor (month 18).
- D 3.2 LLM-extracted literature dataset with confidence information and documented source and processing history (month 24).
- M 3.2 Factors requiring human extraction or review identified before integration into final modelling (month 24).

8.4 Work Package 4 - Primary-Data Enrichment and Expert Validation

Start: Month 12 **End:** Month 30

Goals

- Adapt Ostrom-style questionnaire material to the 24-factor framework and its mapping rules.
- Use AIIS and partner expertise to validate selected variables and data gaps.
- Integrate primary-data enrichment as a secondary validation layer.

Content

This work package uses the applicant's AIIS cooperation context to test and enrich the framework in contemporary irrigation settings. It will synthesise original Ostrom questionnaire material and later developments into a compact questionnaire aligned with the 24 factors and their operational indicators. Depending on partner availability, selected cases or expert reviews from Asian irrigation systems will be used to validate whether the framework and mapping rules capture relevant social, ecological and institutional variation.

Milestones and Deliverables

- D 4.1 Framework-aligned questionnaire module (month 18).
- M 4.1 Partner validation workshop or structured expert review completed (month 24).
- D 4.2 Enrichment and validation report integrated into CoopAt documentation (month 30).

8.5 Work Package 5 - Data Integration, Scientific Modelling and Dissemination

Start: Month 18 **End:** Month 36

Goals

- Integrate harmonised, extracted and enriched data into analysis-ready data frames with dimension-specific outcome construction.
- Estimate and compare GLM, random forest, gradient boosting and neural-network models for factors and configurations associated with SES success.
- Produce a scientific synthesis of robust patterns, synergies, tradeoffs and context limits, and disseminate reusable data documentation and publications.

Content

The final work package builds the analysis-ready CoopAt dataset and uses it to model dimensions of SES success. The team will run preprocessing, missing-data handling, model comparison and robustness checks by outcome dimension, model family, resource domain and region. The scientific synthesis will report which factors and configurations are robustly associated with specific outcome dimensions, where synergies and tradeoffs appear, and where findings remain domain-specific. Results will be disseminated through peer-reviewed articles, conference presentations, visualisations and reusable documentation that allows other researchers to inspect the framework and reproduce the analysis.

Milestones and Deliverables

- D 5.1 Integrated CoopAt analysis dataset with transformation and outcome-construction documentation (month 24).
- D 5.2 Scientific findings and model-robustness report on factors, configurations and tradeoffs in multidimensional SES success (month 30).
- D 5.3 Visualisation, API documentation and dissemination package (month 36).
- M 5.1 Final scientific synthesis completed (month 36).
- M 5.2 Replication documentation, description of source and processing history and code release prepared by month 36.

9 Qualifications

The project leader and principal investigator (PI), Ulrich Frey, is Full Professor at the Institute of Environmental Systems Sciences at the University of Graz, where he holds the Chair of Systems Sciences within the Faculty of Environmental, Regional and Educational Sciences. His research profile is a direct match for CoopAt: he combines substantive commons and cooperation research with quantitative, large-scale comparative analysis of institutions and sustainability outcomes.

On the substantive side, he is the author of the monograph *Success in the Commons: Actors, Institutions and Sustainability* (Oxford University Press, 2020) and of a German-language Springer volume on the sustainable management of natural resources (2018). His empirical and methodological papers, including work on self-organisation in the commons and on defining success in commons governance, engage directly with the Ostrom tradition, design principles and the operationalisation of cooperation outcomes. He has therefore spent more than a decade working on the conceptual problem CoopAt addresses at scale: how to compare, code and learn from heterogeneous SES cases without losing analytical rigour.

On the methodological side, the PI has a strong record of building, validating and publishing quantitative models of complex SES and energy systems, with publications in *Nature Communications*, *Ecology and Society*, *Environmental Modelling & Software* and the *Journal of Artificial Societies and Social Simulation*. This includes long-standing co-development and use of open-source research software, with hands-on experience in pipeline design, validation against expert benchmarks, reproducibility and open release under Findable, Accessible, Interoperable and Reusable (FAIR) principles. These competences are directly relevant to the planned database mapping work, LLM-assisted extraction workflow, integrated data frames and machine-learning models.

The PI also brings the project-leadership experience expected for a three-year FWF project. Prior to his appointment in Graz he led modelling work and third-party-funded projects in the Department of Systems Analysis and Technology Assessment at the German Aerospace Center (DLR), supervised doctoral and postdoctoral researchers, and contributed to the infrastructure of an internationally visible research group. At the University of Graz he is anchored in the *Complexity of Life* research profile area, which provides an interdisciplinary network in environmental systems sciences, sustainability transitions, computational social science and data-driven research. His international network in commons scholarship, including IASC and the Ostrom Workshop community, directly supports the cooperation and validation arrangements described in Section 7.

Project Team

The project will be carried out by the PI together with two predoctoral researchers and cooperation input from AIIS partners. One predoctoral position will focus on computational methods: document processing, LLM-assisted extraction, validation workflows, data integration and machine-learning models. The second predoctoral position will focus on SES and commons analysis: framework refinement, database harmonisation, codebook development, expert validation and interpretation of model results. The two positions are deliberately complementary but overlapping, so that core project knowledge is shared and the team remains robust if individual tasks take longer than expected.

Both positions are suitable for doctoral theses. The computational thesis can address reliable extraction and modelling of SES case evidence from unstructured text. The SES-oriented thesis can address the comparative operationalisation of social, ecological and institutional factors across resource-governance cases. Recruitment will follow open, internationally advertised calls in line with University of Graz policy.

Career Development and Supervision

Both predoctoral researchers will receive structured supervision by the PI, with regular individual meetings, written progress reviews and a yearly thesis advisory committee in line with faculty practice. They will be embedded in the Doctoral Academy of the University of Graz and in the interdisciplinary environment of Environmental Systems Sciences. The project will support presentations at venues such as the IASC Biennial Conference, sustainability and environmental-systems conferences, and selected computational social-science or natural language processing venues where the extraction workflow is methodologically relevant.

The work plan gives both early-career researchers publishable responsibilities. Expected outputs include a framework and harmonisation paper, an LLM-assisted extraction and validation paper, and one or more modelling papers on the determinants of sustainable SES governance. Where partner access permits, the AIIS enrichment component will also provide opportunities for international research exchange and applied collaboration.

Infrastructure and Institutional Embedding

The University of Graz and the Institute of Environmental Systems Sciences provide office space, IT equipment, software licences, access to bibliographic databases, broadband connectivity, library services, and administrative, legal and research-service support. Workstations with graphics processing units (GPUs) at the institute cover routine LLM inference, prompt development and local model deployment, while the university's high-performance computing facilities are available for larger extraction and modelling runs. The institutional setting in Graz, combined with the PI's international commons network and AIIS cooperation context, gives CoopAt the substantive expertise, computational infrastructure and dissemination channels needed to deliver its results.

10 Ethical Aspects

We commit to adhere to the rules of sound scientific practice as established by the Austrian Agency for Research Integrity [31], the principles and good research practices established by the All European Academies [32], as well as *Ethics for Researchers* [33].

The main empirical material for CoopAt consists of already-published academic case studies, grey-literature reports and existing SES databases. The project does not involve experiments on humans or animals, clinical interventions, or collection of sensitive personal data. The targeted primary-data enrichment component is limited to framework validation, expert input and selected case information obtained through cooperation partners. If any activity moves beyond expert consultation and uses interviews or questionnaires with identifiable persons, the team will seek the required institutional ethics review before data collection begins.

Three ethical aspects are particularly relevant.

Handling of Case-Study Material and Affected Communities

Many SES case studies describe identifiable communities, such as irrigation associations, forest user groups, fishery organisations or local resource-governance institutions. Although much of this information is already available in the academic record, CoopAt reaggregates it at a scale that did not previously exist. The project will therefore avoid exposing information in ways that could identify vulnerable groups or invite harmful interpretation.

We mitigate this risk in three ways:

- results will be reported primarily at the level of structural factors, resource domains and modelled patterns rather than as rankings of named communities
- each structured value will retain a source reference and, where possible, an evidence passage or comment so that users can interpret the data in context
- where case material concerns indigenous or otherwise vulnerable communities, the project will consider the Collective benefit, Authority to control, Responsibility and Ethics (CARE) principles for Indigenous Data Governance [34] alongside FAIR data principles.

The Data Ethics Canvas of the Open Data Institute [35] will be used as a structured checklist when new data sources or partner-supplied materials are integrated.

Artificial Intelligence Ethics and European Union Artificial Intelligence Act Compliance

The project uses LLMs for information extraction and machine-learning models for analysis. These methods raise risks of hallucination, bias, opacity and weak reproducibility. CoopAt addresses these risks through documentation and human oversight:

- extraction prompts, model versions, codebook definitions and pipeline configurations will be documented
- a validation sample will be used to measure extraction quality by factor
- low-confidence or contradictory outputs will be routed to human review
- results will distinguish expert-generated, LLM-extracted and human-reviewed values
- open-source models will be used where feasible to support reproducibility and reduce dependence on a single provider.

The use of LLMs in this project is a research-synthesis application and does not deploy artificial intelligence (AI) systems for high-risk use cases such as employment, education, law enforcement, biometric identification or critical infrastructure management. The project will nevertheless align documentation, transparency and human-oversight practices with the European Union (EU) Artificial Intelligence Act [36].

Responsible Modelling and Dissemination

Predictive models of social-ecological governance can be misinterpreted if presented as deterministic recommendations. CoopAt will therefore report model uncertainty, domain limits and robustness checks explicitly. Visualisations and publications will be framed around improving understanding of sustainable governance, not around scoring communities or prescribing one-size-fits-all institutional designs. Together with the gender-related considerations described in Section 11, these measures cover the ethical obligations currently foreseen for the project.

11 Gender Aspects

Gender is relevant to CoopAt both in team composition and in research content. For recruitment, the University of Graz procedures for open calls, transparent selection criteria and equal-opportunity hiring will be followed. The project will aim for gender-balanced candidate pools for the two predoctoral positions and will use structured evaluation criteria to reduce implicit bias. Flexible working arrangements and family-friendly supervision practices will be offered in line with university policy.

Gender is also a substantive dimension of SES. In many irrigation, forestry and fishery contexts, access to resources, participation in decision making, labour contributions, monitoring responsibilities and exposure to environmental risk differ by gender. If case studies underreport women's roles or treat households as homogeneous units, a database built from those studies can reproduce that bias. The project will therefore include gender-related information in the codebook where the source material permits and will record whether gender dimensions are absent, explicit or ambiguous in each case.

During modelling and interpretation, the team will treat gender variables carefully. Where data coverage is sufficient, gender composition, participation and representation will be analysed as possible factors in governance outcomes. Where coverage is weak, the project will report this as a limitation rather than drawing unsupported conclusions. This approach makes gender both a quality-control issue for extraction and a substantive research dimension for understanding sustainable SES governance.

A Annex 1 - List of references

[1] Yixu Wang, Jie Gong, Zhihui Yang, and Yuehua Zhu. Social-ecological system research in a changing world: State of the art and future challenges. *Journal of Cleaner Production*, 489:144725, 2025. doi: 10.1016/j.jclepro.2025.144725.

[2] Claudia R. Binder, Jochen Hinkel, Pieter W. G. Bots, and Claudia Pahl-Wostl. Comparison of frameworks for analyzing social-ecological systems. *Ecology and Society*, 18(4):26, 2013. doi: 10.5751/ES-05551-180426. URL `https://www.ecologyandsociety.org/vol18/iss4/art26/`.

[3] Amy R. Poteete and Elinor Ostrom. Fifteen years of empirical research on collective action in natural resource management: Struggling to build large-n databases based on qualitative research. *World Development*, 36(1):176–195, 2008. doi: 10.1016/j.worlddev.2007.02.012.

[4] Maja Schlüter, Reinette Biggs, Hayley Clements, Alta de Vos, Kristine Maciejewski, and Rika Preiser. Synthesis and emerging frontiers in social-ecological systems research methods. In Reinette Biggs, Alta de Vos, Rika Preiser, Hayley Clements, Kristine Maciejewski, and Maja Schlüter, editors, *The Routledge Handbook of Research Methods for Social-Ecological Systems*, pages 453–479. Routledge, London, 2021. doi: 10.4324/9781003021339-40.

[5] Ulrich J. Frey. Multiple method modelling reveals lack of robustness in natural resource management research. *Journal of Environmental Management*, 281:111812, 2021. doi: 10.1016/j.jenvman.2020.111812.

[6] Jacopo A. Baggio, Allain J. Barnett, Irene Perez-Ibarra, Ute Brady, Elicia Ratajczyk, Nathan Rollins, Cathy Rubiños, Hoon C. Shin, David J. Yu, Rimjhim Aggarwal, John M. Anderies, and Marco A. Janssen. Explaining success and failure in the commons: The configural nature of ostrom's institutional design principles. *International Journal of the Commons*, 10(2):417–439, 2016. doi: 10.18352/ijc.634. URL `https://www.thecommonsjournal.org/article/10.18352/ijc.634/`.

[7] Michael Cox, Georgina G. Gurney, John M. Anderies, Eric Coleman, Emily Darling, Graham Epstein, Ulrich J. Frey, Mateja Nenadovic, Edella Schlager, and Sergio Villamayor-Tomas. Lessons learned from synthetic research projects based on the ostrom workshop frameworks. *Ecology and Society*, 26(1):art17, 2021. doi: 10.5751/ES-12092-260117. URL `https://www.ecologyandsociety.org/vol26/iss1/art17/`.

[8] John Dagdelen, Alexander Dunn, Sanghoon Lee, Nicholas Walker, Andrew S. Rosen, Gerbrand Ceder, Kristin A. Persson, and Anubhav Jain. Structured information extraction from scientific text with large language models. *Nature Communications*, 15(1):1418, 2024. doi: 10.1038/s41467-024-45563-x.

[9] Allain J. Barnett, Stefan Partelow, Ulrich Frey, Alejandro García-Lozano, Maria Del Mar Mancha-Cisneros, Christoph Oberlack, Elicia Ratajczyk, Hillary Smith, Sergio Villamayor-Tomás, and Charlotte K. Whitney. Defining success in the commons. *International Journal of the Commons*, 14(1):366–387, 2020. doi: 10.5334/ijc.994.

[10] Claudia Pahl-Wostl, Xavier Basurto, and Sergio Villamayor-Tomas. Comparative case study analysis. In Reinette Biggs, Alta de Vos, Rika Preiser, Hayley Clements, Kristine Maciejewski, and Maja Schlüter, editors, *The Routledge Handbook of Research Methods for Social-Ecological Systems*, pages 282–294. Routledge, London, 2021. doi: 10.4324/9781003021339-24.

[11] Elinor Ostrom. A general framework for analyzing sustainability of social-ecological systems. *Science*, 325(5939):419–422, 2009. doi: 10.1126/science.1172133.

[12] Elinor Ostrom. *Understanding Institutional Diversity*. Princeton University Press, Princeton, NJ, 2005. ISBN 978-0-691-12207-6 978-1-4008-3173-9. doi: 10.1515/9781400831739.

[13] Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, 1990. ISBN 978-0-521-37101-8 978-0-521-40599-7 978-0-511-80776-3. doi: 10.1017/CBO9780511807763. URL `https://www.cambridge.org/core/product/identifier/9780511807763/type/book`.

[14] Amy R. Poteete and Elinor Ostrom. Heterogeneity, group size and collective action: The role of institutions in forest management. *Development and Change*, 35(3):435–461, 2004. doi: 10.1111/j.1467-7660.2004.00360.x.

[15] Ulrich J. Frey. A synthesis of key factors for sustainability in social-ecological systems. *Sustainability Science*, 12(4):507–519, 2017. doi: 10.1007/s11625-016-0395-z.

[16] Michael Cox, Gwen Arnold, and Sergio Villamayor Tomás. A review of design principles for community-based natural resource management. *Ecology and Society*, 15 (4):38, 2010.

[17] Amy R. Poteete, Marco A. Janssen, and Elinor Ostrom. *Working Together: Collective Action, the Commons, and Multiple Methods in Practice*. Princeton University Press, Princeton, 2010.

[18] Michael J. Bernstein, Maria Del Mar Mancha-Cisneros, Madeline Tyson, Ute Brady, Cathy Alida Rubiños, Hoon Cheol Shin, Sechindra Vallury, Skaidra Smith-Heisters, and Elicia Ratajczyk. Mapping ostrom's common-pool resource systems coding handbook to the coupled infrastructure systems framework to enable comparative research. *International Journal of the Commons*, 13(1):528–552, 2019. doi: 10.18352/ijc.904.

[19] Elicia Ratajczyk, Ute Brady, Jacopo A. Baggio, Allain J. Barnett, Irene Perez-Ibarra, Nathan Rollins, Cathy Rubiños, Hoon C. Shin, David J. Yu, Rimjhim Aggarwal, John M. Anderies, and Marco A. Janssen. Challenges and opportunities in coding the commons: Problems, procedures, and potential solutions in large-n comparative case studies. *International Journal of the Commons*, 2016. doi: 10.18352/ijc.652. URL `https://www.thecommonsjournal.org/article/10.18352/ijc.652/`.

[20] Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, Yang Wang, and Enhong Chen. Large language models for generative information extraction: A survey. *Frontiers of Computer Science*, 18(6):

186357, 2024. doi: 10.1007/s11704-024-40555-y. URL https://link.springer.com/10.1007/s11704-024-40555-y.

[21] OpenAI. GPT-4 technical report, 2024. URL http://arxiv.org/abs/2303.08774.

[22] Gemini Team. Gemini: A family of highly capable multimodal models, 2025. URL http://arxiv.org/abs/2312.11805.

[23] Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiachong Feng, Lingpeng Kong, and Qi Liu. Multimodal ArXiv: A dataset for improving scientific comprehension of large vision-language models, 2024. URL http://arxiv.org/abs/2403.00231.

[24] Zackary Okun Dunivin. Scalable qualitative coding with LLMs: Chain-of-thought reasoning matches human performance in some hermeneutic tasks, 2025.

[25] Zackary Okun Dunivin. Scaling hermeneutics: A guide to qualitative coding with LLMs for reflexive content analysis. *EPJ Data Science*, 2025. doi: 10.1140/epjds/s13688-025-00548-8. URL https://epjdatascience.springeropen.com/articles/10.1140/epjds/s13688-025-00548-8.

[26] Hao Lin and Yongjun Zhang. Navigating the risks of using large language models for text annotation in social science research. *Social Science Computer Review*, 44(3):403–427, 2026. doi: 10.1177/08944393251366243. URL https://journals.sagepub.com/doi/10.1177/08944393251366243.

[27] Petter Törnberg. Best practices for text annotation with large language models. *Sociologica*, 18(2):67–85, 2024. doi: 10.6092/ISSN.1971-8853/19461. URL https://sociologica.unibo.it/article/view/19461.

[28] Ben Nagel and Stefan Partelow. A methodological guide for applying the social-ecological system (SES) framework: A review of quantitative approaches. *Ecology and Society*, 27(4):art39, 2022. doi: 10.5751/ES-13493-270439. URL https://www.ecologyandsociety.org/vol27/iss4/art39/.

[29] Ulrich J. Frey and Hannes Rusch. Using artificial neural networks for the analysis of social-ecological systems. *Ecology and Society*, 18(2):40, 2013. doi: 10.5751/ES-05202-180240.

[30] Ulrich J. Frey and Hannes Rusch. Modeling ecological success of common pool resource systems using large datasets. *World Development*, 59:93–103, 2014. doi: 10.1016/j.worlddev.2014.01.034.

[31] Austrian Agency for Research Integrity (OeAWI). Guidelines for good scientific practice. https://oeawi.at/en/guidelines/, 2015.

[32] ALL European Academies (ALLEA). The european code of conduct for research integrity, revised edition. https://allea.org/code-of-conduct/, 2017.

[33] European Commission. Ethics for researchers: Facilitating research excellence in fp7, 2013.

[34] Stephanie Russo Carroll, Ibrahim Garba, Oscar L. Figueroa-Rodríguez, Jarita Holbrook, Raymond Lovett, Simeon Materechera, Mark Parsons, Kay Raseroka, Desi Rodriguez-Lonebear, Robyn Rowe, Rodrigo Sara, Jennifer D. Walker, Jane Anderson, and Maui Hudson. The CARE principles for indigenous data governance. *Data Science Journal*, 19:43, 2020. doi: 10.5334/dsj-2020-043. URL `http://datascience.codata.org/articles/10.5334/dsj-2020-043/`.

[35] Open Data Institute. The data ethics canvas. `https://theodi.org/insights/tools/the-data-ethics-canvas-2021/`, 2021.

[36] European Parliament and Council. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L series, 12 July 2024, 2024. URL `http://data.europa.eu/eli/reg/2024/1689/oj`.

B Annex 2 - Financial aspects

a) Details on the research institution of the applicant

Existing personnel

Ulrich Frey is already employed full-time by the University of Graz. His work on project leadership, supervision, conceptual integration and dissemination will therefore be provided as an in-kind contribution and is not requested from the FWF budget. The project team will also receive support from administrative personnel, the research service department, the library and the IT department of the university.

Existing infrastructure

The University of Graz provides office space, standard computers, software licences, access to bibliographic databases, broadband connectivity, library services and administrative support. The Institute of Environmental Systems Sciences has workstations suitable for routine data processing, prompt development and LLM-assisted extraction experiments. University high-performance computing facilities are available for larger extraction and modelling runs. No major equipment purchase is requested at this stage.

b) Information on the funding requested

The requested budget is a provisional, internally consistent calculation based on the FWF 2026 personnel rates. It will be aligned with the final University of Graz cost breakdown before submission. Direct project costs are estimated at € 380,508. The 5% general project costs amount to € 19,025. The provisional total requested funding is therefore € 399,533.

Personnel costs: € 314,308

The central cost item is personnel for two doctoral candidates at 30 hours per week over the 36-month project duration. The 2026 FWF doctoral candidate rate is € 50,220 per year for 30 hours per week. Applying the FWF annual increase of 4.25% from the second project year onward gives € 50,220 in year 1, € 52,354.35 in year 2 and € 54,579.41 in year 3. This results in € 157,153.76 per doctoral candidate and € 314,307.52 for two positions.

One doctoral candidate will cover computational tasks: document processing, LLM-assisted extraction, validation workflows, data integration and machine-learning analysis. The second doctoral candidate will cover SES and commons tasks: framework refinement, database harmonisation, codebook development, expert validation and interpretation of model results. Both positions are scientifically necessary because the project depends on close integration of computational methods and substantive SES expertise.

Equipment costs: € 0

No dedicated equipment is requested. Existing institute workstations and university high-performance computing infrastructure are expected to cover routine computational needs. If a specific repair or small upgrade becomes necessary, it will be handled only if it is project-specific and eligible under the applicable FWF rules.

Material, services and dissemination: € 32,000

Material and service costs cover the project-specific costs needed to make extraction, validation and dissemination reproducible. The provisional allocation is € 18,000 for LLM application programming interface (API) use and benchmarking against open-source models, € 6,000 for document processing, optical character recognition (OCR), data cleaning and technical services, € 4,000 for research-data management and long-term archiving, and € 4,000 for dissemination materials, visualisation preparation and project documentation. Open-access publication fees are not requested here because FWF open-access support is handled through the relevant publication funding routes.

Travel, cooperation and primary-data enrichment: € 34,200

Travel and field-related costs are included because AIIS cooperation and expert validation are part of the validation strategy. The provisional allocation is € 25,000 for AIIS-related partner coordination, validation activities and selected questionnaire-based enrichment, € 6,000 for three international conference participations, and € 3,200 for four national or European workshops and conferences. The primary-data enrichment component remains secondary to the core framework, harmonisation, extraction and modelling work, so delays in field access do not endanger project completion.

General project costs: € 19,025

General project costs are calculated as 5% of the direct project costs. They cover small, project-related expenses that cannot be assigned reliably to a specific direct cost line at the proposal stage. They are not used to inflate personnel, equipment or publication-cost requests.

C Annex 3 - Curricula vitae (CVs) and description of previous research achievements

Univ.-Prof. Dr. Dr. Ulrich Frey

Ulrich Frey is Full Professor at the Institute of Environmental Systems Sciences at the University of Graz, where he holds the Chair of Systems Sciences within the Faculty of Environmental, Regional and Educational Sciences. His research combines commons and cooperation research, SES, quantitative modelling and reproducible computational workflows. This profile directly matches CoopAt, whose central task is to turn heterogeneous SES case evidence into a validated comparative data and modelling infrastructure.

Previous Research Achievements

The applicant has worked for more than a decade on the conceptual and empirical problem at the centre of this proposal: how social, ecological and institutional factors can be compared across heterogeneous common-pool resource systems without losing analytical rigour. His monograph *Success in the Commons: Actors, Institutions and Sustainability* (Oxford University Press, 2020) and his German-language Springer volume on sustainable natural-resource management develop this research line in depth. His empirical and methodological publications address self-organisation in the commons, the operationalisation of sustainability outcomes, and quantitative modelling of complex SES cases.

The applicant also has a strong record in computational modelling and research software. His work has appeared in outlets including *Nature Communications*, *Ecology and Society*, *Environmental Modelling & Software* and the *Journal of Artificial Societies and Social Simulation*. This includes experience with open-source pipelines, validation against expert benchmarks, reproducibility, FAIR data principles and machine-learning models for complex SES and energy systems. These competences are directly relevant to CoopAt's database mapping work, LLM-assisted extraction workflow, integrated data frames and model-comparison strategy.

Project Leadership and Networks

Before joining the University of Graz, the applicant led modelling work and third-party-funded projects in the Department of Systems Analysis and Technology Assessment at the German Aerospace Center (DLR). He has supervised doctoral and postdoctoral researchers and contributed to the infrastructure of an internationally visible research group. At the University of Graz, he is embedded in Environmental Systems Sciences and the *Complexity of Life* research profile area, which provides an interdisciplinary environment for sustainability, systems science, computational social science and data-driven environmental research.

The applicant's international commons network, including the International Association for the Study of the Commons and the Ostrom Workshop community, supports expert review, data-source identification and dissemination. His involvement in AIIS provides a concrete cooperation context for validating the 24-factor framework and enriching selected irrigation cases without making fieldwork the dominant dependency of the project.

Career Development Contribution

CoopAt creates two complementary doctoral projects. One doctoral candidate will develop and evaluate reliable extraction and modelling of SES case evidence from unstructured text. The second will work on comparative operationalisation of social, ecological and institutional factors across resource-governance cases. Both researchers will receive structured supervision, publishable responsibilities and access to the applicant's interdisciplinary and international research networks.

D Annex 4 - Statement on environmental sustainability in the research process

CoopAt is a data-intensive but comparatively low-material research project. Its main environmental impacts arise from computing, travel and ordinary office work rather than from laboratory equipment, consumables or experimental materials. The project will therefore focus on those aspects of ecological sustainability that are directly within the research team's sphere of influence.

Travel and Meetings

Project coordination will use virtual and hybrid meetings by default where they do not reduce scientific quality. International travel will be concentrated on activities that require in-person participation, such as selected conferences, AIIS partner coordination and validation workshops. Within Austria and neighbouring countries, train travel will be preferred where feasible. Multiple project purposes will be combined into single trips where possible, for example by linking partner meetings with workshops or conference participation.

Computing and Data Processing

The project will reuse existing University of Graz workstations and high-performance computing infrastructure instead of purchasing dedicated hardware. LLM use will be staged: small pilot batches, prompt testing and open-source model trials will precede larger extraction runs. This reduces unnecessary compute consumption and avoids repeated large-scale processing of poorly specified tasks. Intermediate outputs, source metadata and validation results will be stored systematically so that extraction runs do not have to be repeated without scientific reason.

Procurement and Dissemination

No major equipment purchase and no material-intensive field campaign are planned. Project documentation, code, visualisations and dissemination material will be distributed digitally wherever possible. Printed material will be limited to situations where it is necessary for workshops, teaching or stakeholder communication. No additional FWF funding is requested specifically for sustainability measures. The planned practices are integrated into the normal research workflow.