Check for updates

RESEARCH ARTICLE

REVISED The European Peatland Policy Portal: Demonstrating Machine Learning and Analytics Tools for Evidence-based Land Use Policymaking

[version 2; peer review: 1 approved, 1 approved with reservations]

M.A. Waskow[1,2], Niall Ó Brolcháin[1,2], Muhammad Yasar Khan[1,2], Yifan Wang[1,2], Waqas Shoukat[1,2], Fergus O'Donoghue[1,2], Saeed Alsamhi[1,2], Sandra Riordan[1,2], Letícia Jurema[3], Fiona Lefebvre[1,4], Ella Cheney[1,5], Terry R. Morley[6]

[1]ASPECT Unit, Insight Research Ireland Centre for Data Analytics, Galway, County Galway, Ireland
[2]Data Science Institute, University of Galway, Galway, County Galway, Ireland
[3]Naturschutzbund Deutschland eV, Berlin, Berlin, Germany
[4]École des Mines d'Albi-Carmaux, Albi, Occitanie, France
[5]University College Dublin, Dublin, County Dublin, Ireland
[6]Discipline of Geography and The Ryan Institute, University of Galway, Galway, County Galway, Ireland

v2 **First published:** 24 Apr 2026, **6**:134
https://doi.org/10.12688/openreseurope.23501.1
Latest published: 29 Jul 2026, **6**:134
https://doi.org/10.12688/openreseurope.23501.2

Abstract

Achieving the ambitions of the European Green Deal requires overcoming longstanding fragmentation in the policy frameworks that govern peatlands. The EU LIFE Multi Peat project introduces the European Peatland Policy Portal as a novel policy dataset and toolkit to support analysis and evidence-based policy development in this disjointed landscape of peatland policy. The portal presents the first comprehensive European Peatland Policy Dataset, with over 500 peatland-affecting policies catalogued across 28 countries and 22 languages, spanning the United Nations, European Union, national, and sub-national levels. Alongside the unified, multilingual policy catalogue, the portal introduces an interactive mapping tool designed to support spatial policy analysis, as well as a retrieval-augmented generation (RAG)-based question-answering tool that allows users to query the full policy dataset with the power of machine learning. With multilingual comprehension, the RAG tool enables users to interface with policies which would otherwise sit behind a language barrier. In its drive to support systematic policy examination and analysis across countries, governance levels, domains, and languages, the portal aims to bridge efficiency gaps and aid the coordination of European

Open Peer Review

Approval Status ✓ ?

	1	2
version 2 (revision) 29 Jul 2026	✓ view	
version 1 24 Apr 2026	? view	? view

1. **Subhankar Das**, Duy Tan University, Da Nang, Vietnam

2. **Mohammad Fazle Rabbi**, University of Debrecen, Debrecen, Hungary

Any reports and responses or comments on the article can be found at the end of the article.

peatland governance. While the portal proposes these tools for peatland policy specifically, the authors argue that the framework is transferable to other land use policy domains that face complex regulatory environments.

Keywords
Environmental policy, evidence-based policy, policy visualisation, policy analysis tools, peatlands, NLP

Corresponding author: M.A. Waskow (margaret.waskow@universityofgalway.ie)

Author roles: Waskow MA: Software, Writing – Original Draft Preparation; **Ó Brolcháin N**: Conceptualization, Funding Acquisition, Writing – Original Draft Preparation; **Khan MY**: Software, Writing – Original Draft Preparation; **Wang Y**: Data Curation, Writing – Original Draft Preparation; **Shoukat W**: Software; **O'Donoghue F**: Data Curation; **Alsamhi S**: Visualization, Writing – Review & Editing; **Riordan S**: Visualization, Writing – Review & Editing; **Jurema L**: Funding Acquisition, Project Administration, Writing – Original Draft Preparation; **Lefebvre F**: Software; **Cheney E**: Data Curation; **Morley TR**: Funding Acquisition, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: This project has received funding from the European Union's LIFE programme under grant agreement No LIFE20 CCM/DE/001802 (LIFE Multi Peat). This project also received funding from Research Ireland as part of Grant Number SFI/12/RC/2289 P2 Insight 2, the Insight Research Ireland Centre for Data Analytics.
The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2026 Waskow MA *et al*. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Waskow MA, Ó Brolcháin N, Khan MY *et al.* **The European Peatland Policy Portal: Demonstrating Machine Learning and Analytics Tools for Evidence-based Land Use Policymaking [version 2; peer review: 1 approved, 1 approved with reservations]** Open Research Europe 2026, **6**:134 https://doi.org/10.12688/openreseurope.23501.2

First published: 24 Apr 2026, **6**:134 https://doi.org/10.12688/openreseurope.23501.1

REVISED **Amendments from Version 1**

The Peatland Policy Dataset and its mentions and analyses have been updated to reflect its new composition, as it has increased to over 500 policies at the time of writing.

Figures 3 and 4 were updated to reflect the updated policy dataset and Figures 5-10 were updated to reflect the design changes to the portal website.

Additional documentation has been added to the European PPP Dataset section concerning the acquisition of policies, namely: constraints of the project on methodology, a table of inclusion and exclusion criteria, a table of worked edge cases, and a comparative coverage analysis of the portal and Overton (additionally a paragraph on Overton has been added to the Related Work section).

Additional documentation has been added to the RAG Question-Answering Tool section, including: information on the system architecture, a quantitative evaluation of the system via the RAGAS framework, and a quantitative comparison of the RAG tool against a non-RAG LLM. Additional references have been added as well as footnotes.

Phrasing throughout the paper concerning the efficacy of tools has been updated so that design intentions are not represented as demonstrated outcomes.

Conclusions about peatland policy drawn from the Policy Collection have been removed to instead analyse the composition of the collection as constructed.

In the Policy Map section, shortcomings of the peatland polygons and policy polygons have been articulated.

In the second sentence, the calculated values of EU peatland coverage have been replaced with the verbatim European region values from the source.

A new version of The European Peatland Policy Dataset (Version July 2026) was added.

A new version of the RAG Evaluation (Quantitative Evaluations of the PPP RAG Tool (Version v3)) Dataset was added.

A new supplementary dataset was added: Partner and Stakeholder Contributions to the PPP's Peatland Policy Collection.

Any further responses from the reviewers can be found at the end of the article

Introduction

Peatlands are critical to Europe's climate and biodiversity strategy, and in particular the EU Green Deal. These terrestrial wetland ecosystems only cover an estimated 5.4% (593,727 km^2) of Europe's surface area, yet they store approximately a third of the region's soil carbon. [1] Despite their importance as natural carbon sinks, decades of drainage for agriculture and industrial use have transformed the majority of European peatland systems into net greenhouse gas (GHG) emitters. Mitigating climate change and achieving the EU's 2050 climate-neutrality target rely on the dual strategy of peatland protection and restoration,[2] a process that requires decisive policy action to bridge the gap between scientific goals and implementation.[2–4]

With the final approval of the EU Nature Restoration Law in June 2024, Member States are now legally mandated to restore 30% of degraded habitats by 2030, including specific targets for rewetting organic soils like peatlands. As of 2026, a primary barrier to achieving these targets is a persistent "data gap"-- policymakers often lack standardised, real-time evidence to monitor peatland condition, verify restoration progress, and quantify GHG mitigation at a landscape scale.[5] Traditional monitoring is often hindered by the remote and inaccessible nature of these terrains, making fieldwork costly and infrequent.

Despite its importance for effective implementation, a review of eight national strategies across Europe found that all were failing to reach the level of policy integration required for meaningful action.[6] Chen et al. (2023) also found that EU peatland governance landscape is currently full of policy conflicts, both vertically across governance levels and horizontally across sectors.[4] Peatland experts have provided recommendations for policies to protect peatlands,[7,8] but evidence-based policy often requires compromise with other factors impacting the selection of policy actions.[9] As peatland policies are scattered across levels and domains, it is difficult for policymakers to efficiently traverse the policy landscape and determine best practice when developing policy solutions in their own contexts.

The European Peatland Policy Portal

With the aim of addressing the diffused nature of peatland policies and tools, this paper introduces the European Peatland Policy Portal (PPP), a novel digital toolkit developed by the ASPECT Unit at the University of Galway through the EU

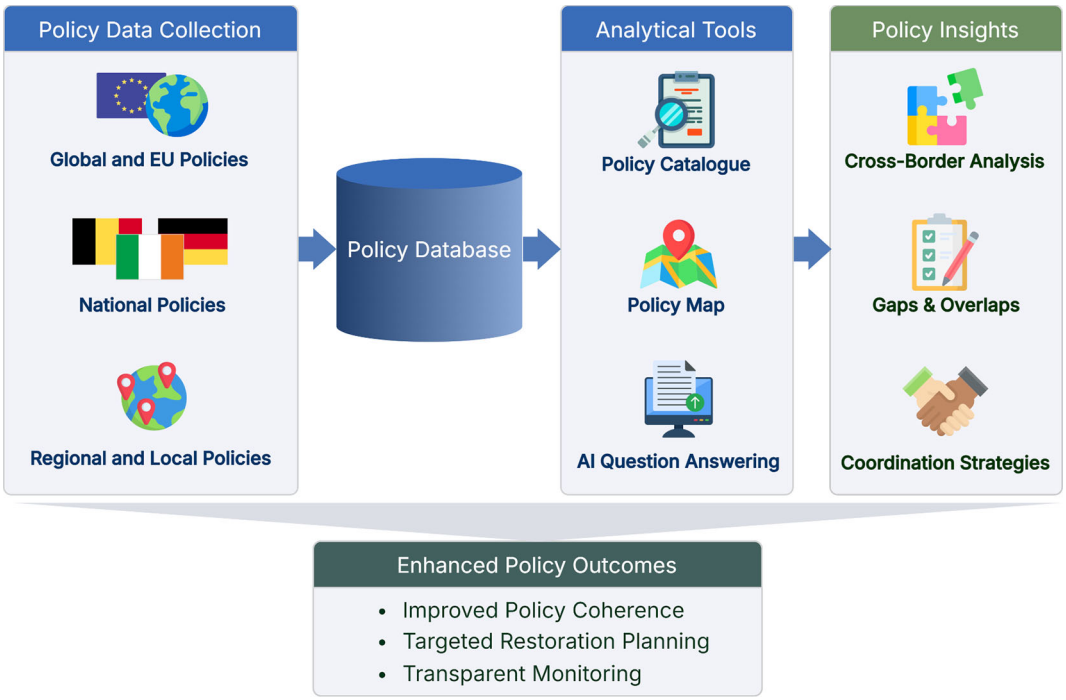


Figure 1. A diagram outlining the contributions of the European Peatland Policy Portal and its intended impacts.

LIFE Multi Peat project together with multiple European partners. The portal is built to serve as a centralised evidence-based policy platform and toolkit that integrates:

- Machine Learning (ML) & AI: A public platform using ML and AI tools to reference over 500 policies from the EU and all EU member states plus the UK.
- Advanced Analytics: Deploying predictive tools, such as the Site Emissions Tool and the Financial Feasibility and Pricing tool to estimate the viability of carbon credit schemes and specialised wetland farming methods for specific sites using the Greenhouse Gas Emissions Site Type (GEST) methodology[10] and a blended financial model.
- Evidence-based Policymaking: Designed to support the analysis, alignment and development of policies by aligning technical data with policies to facilitate cross-sectoral collaboration.

By demonstrating these novel tools, the authors argue that high-resolution, data-driven insights can bridge the gap between scientific observation and legislative implementation, helping to ensure that European land use policies are both effective and verifiable in the face of the climate crisis.

The contributions of this paper to the realm of European public policy are threefold. First, it presents a curated peatland policy dataset as a unified collection of global, EU-level, national, and subnational peatland-affecting policies scattered across disjointed, decentralised regional, national, and international policy repositories. Secondly, it demonstrates an interactive map tool which allows policy researchers the opportunity to engage with the policy landscape on a geospatial level. Finally, it proposes a retrieval-augmented generation (RAG)-based question-answering tool that enables users to engage with the entire policy dataset in a matter of seconds. Through these contributions, the authors hope to facilitate more efficient analysis and development of policies pertaining to peatlands in Europe (Figure 1).

The Project Perspective

The EU LIFE Multi Peat project (2021–2026) is the result of the urgent need to address the climate crisis and contribute to reaching national and global targets of carbon emissions reductions. A consortium of partners from peatland-rich countries --Poland, Germany, Netherlands, Belgium, and Ireland-- united to rewet and monitor degraded peatlands,

promote the exchange of knowledge in key aspects of peatland restoration, and explore the driving forces that impact restoration at a greater scale.

The Multi Peat project developed the European Peatland Policy Portal in the frame of this consortium to contribute to a more cohesive and integrated approach to peatland restoration through policy. Thus, the portal aims to address policy and knowledge gaps as well as fragmented and conflicting legislations within and across European countries, which represent significant obstacles to larger scale restoration, hindering effective climate change mitigation initiatives.

The combined experience within the project consortium and its networks of navigating peatland policies and legislation to implement long-term peatland restoration was essential in creating the portal. Establishing a platform that hosts a Europe-wide peatland policy catalogue and various policy development tools that support policy analysis, including an interactive map of the peat soils of the EU and the UK, the portal is intended to facilitate the development and implementation of integrated approaches, as well as improve the knowledge base for governance of peatlands and restoration initiatives.

Portal Mission And Intended Audience

The portal works with data providers to gather and display data relating to the most important policies affecting the protection, restoration and management of peatlands at regional and transnational levels. Through this, it provides a cross-referencing tool designed to aid evidenced-based policy development. In addition, the platform showcases examples of peatland protection and restoration activities, highlighting enabling factors as well as challenges in implementation.

The portal strives to promote evidence-based policy development and support integrated restoration across regional and transnational levels by engaging a wide array of stakeholders. These range from direct decision-makers, such as politicians and public servants, to indirectly involved actors like researchers, journalists, and lobby groups. To maximise utility for these key users, the platform maintains a commitment to regular updates.

The portal is designed to help stakeholders understand regional and national differences in peatland policies, identify best practices of policy development, and obtain ideas on how to improve policies. Additionally, the portal aims to address actors that seek to develop or engage in different peatland projects, including non-governmental organisations (NGOs), government bodies, private landowners or fund holders, by presenting, among others, best practices, funding opportunities, and actor networks. Further, the portal provides tools intended to help users estimate the financial viability of peatland restoration initiatives and calculate GHG emissions before and after rewetting scenarios.

Related Resources and Work

Policy Portals

The portal's Policy Collection draws inspiration from EUR-Lex [1] and Carbon Gap's Policy Tracker [2], which both have created centralised repositories of policies from scattered sources. EUR-Lex, created and maintained by the Publications Office of the European Union, is a multilingual repository of EU legal documents which presents users with the ability to explore not only the documents, but their metadata and connections to other documents as well. While the site is mostly focused on international law, it does include the texts of national transposition measures for some member states, and includes references to both national transposition measures and national case-law related to EU law. This repository is a complete, regularly-updated record of EU-level legal documents, but national and subnational policies are out of its scope.

Carbon Gap's Policy Tracker, with a focus on policy topic rather than level, aims to enable the systematic assessment and comparison of carbon dioxide removal (CDR) policy environments in different countries around the world. The tracker provides a collection of national policies concerning CDR, country-specific profiles of the CDR policy landscape, and the site's own determinations of best practice in the area. Like the PPP, this resource records national policies, but it does not include sub-national policies; there are also minor overlaps in the CDR and PPP policy datasets since healthy peatlands' carbon sequestration capacities make peatland restoration a topic of interest for CDR.

The UK R&D Policy Toolkit [3], though it focuses on research and development policy instead of peatlands, also aligns with the PPP in its mission to provide resources, tools, and evidence for policy development.

[1] https://eur-lex.europa.eu/
[2] https://tracker.carbongap.org/
[3] https://rd-toolkit.britishprogress.org/

Overton [4] is a policy and grey literature database which contains over 25 million documents from 193 countries. Originally designed to help researchers explore the impact of their work on policies and other instances of decision-making, its current subscription-based database spans various kinds of documents from government and non-government sources (e.g. policy documents, reports, transcripts, factsheets, think tank publications, etc.). Overton and the PPP share the same general goal of helping to streamline policy discovery and analysis by centralising difficult-to-find policies and resources in one place, however an analysis described later in the paper reveals that the PPP's specialisation and input from experts has resulted in a policy collection that fills in some of Overton's gaps surrounding peatland policy.

Spatial Policy Analysis

The interactive Peatland Policy Map introduces a new paradigm for applying geospatial analysis in the policy domain. In policy contexts, spatial analysis is most often used to understand data when defining a problem or opportunity,[11,12] to determine optimal policy solutions by predicting the effects of proposed actions,[13,14] or to evaluate policy impacts.[15,16] Reinvestment Fund's Policy Map [5] provides interactive maps and data analysis tools for these sorts of policy-oriented spatial analyses, focused on statistical geospatial data. In contrast, the PPP's Policy Map tool is designed to enable spatial analysis of policies themselves; by assigning each policy a geometry based on its administrative extent, the tool provides spatial visualisation of the policy landscape, showing all the policies across departments and administrative levels that affect a single coordinate.

AI for Policy Analysis

The portal's RAG-based Question-Answering tool follows a growing trend of integrating AI into public policy research through both analytical[17–19] and generative[20–23] machine learning methods, specifically natural language processing (NLP). Policy Bot, of Nagarajan et al. (2025), receives a policy document from its user and provides question-answering in a dialogue format, drawing its responses from the policy document's text using retrieval-augmented generation.[21] While the portal's RAG tool does not operate in a dialogue format, it does generate its responses by drawing from the hundreds of policies in the portal's policy dataset, and its generated text response is accompanied by the policy documents used to create the response, ranked for contribution. This added explainability addresses a key concern of the trustworthiness of incorporating generative AI into public administration.[24]

The European Peatland Policy Dataset

The project first began its attempt to address the fragmented nature of the peatland policy landscape by gathering the contemporary peatland-affecting policies of Europe into one central catalogue. Policy documents were collected according to the following guidelines and were recorded with additional metadata for later filtering and display in the portal's Policy Collection.

Dataset Curation

The discovery and annotation task for building the European peatland policy dataset occurred in three phases. As guided by the Multi Peat project deliverables, the peatland policy collection's construction included not only web and database searches but also collaborations with project partners and stakeholder workshops. While the inclusion of policies suggested by experts introduces selection bias into the policy collection process, translation challenges and the decentralised nature of many national and sub-national policy repositories left gaps in the search method's results which experts of the country's policy were able to fill. Thus, as it stands, the policy collection is not presented as a necessarily comprehensive or exhaustive list of the peatland-affecting policies of Europe, but rather as an ongoing attempt to consolidate these policies based on documented criteria and expert input.

A snapshot of the policy collection at the time of writing (571 policies) has been saved and archived on Zenodo.[25] As the acquisition process took place over several phases, the archived dataset includes the date of each policy entry's creation as a draft into our private collection, as well as the date of the policy entry's most recent publication to the public portal following either the initial completion or subsequent update of its metadata. Additional documentation of input to the policy collection from partners and stakeholders has also been archived on Zenodo, including policies from survey responses and outcomes from workshop discussions.[26]

Phase I (September 2023–September 2024)

In the first phase of policy collection, approximately 100 policies were acquired from workshops with policy researchers in the Multi Peat project's partner countries, as well as found by searching Google for country names and "peatland" or

[4]https://www.overton.io/
[5]https://www.policymap.com/

Table 1. List of the 19 keywords or keyword sets drawn from key peatland-affecting EU laws and used to search for national implementations.

Habitats directive	Soil strategy
Natura 2000	Nature restoration regulation/nature restoration law
Biodiversity strategy	Environmental impact assessment/eia
Forestry strategy	Farm to fork/food system
Common agricultural policy	Land use, land use change, and forestry/lulucf
Birds directive	Carbon removals certification/crcf
Methane strategy	National energy and climate plans/necp
Water framework	Renewable energy/red iii
Taxonomy regulation	Greenhouse gas emission/eu ets
Climate-adapt	

broad categories like "biodiversity". These policies and their annotations were added to a database in the portal's content management system, Odoo [6].

Phase II (October 2024–September 2025)

The second phase of discovery and annotation centred around national adaptations of EU-level directives, regulations, frameworks, strategies, etc. The collection process began by conducting a keyword search—"peatland" and "peat" especially—on EUR-Lex to find documents relevant to peatland conservation and peat or peatland use. After internal discussions and consultation with subject experts, the process identified 19 key EU laws and strategies which became the 19 sets of keywords used in the rest of the second phase, as can be seen in Table 1.

In the continuation of the second phase, policy researchers used the keywords and their linguistic equivalents to search for relevant national or subnational policies in EU member states and the UK. After first identifying a document from the keyword search, the next step was to determine its eligibility to be added to the dataset.

In order to obtain the most explicitly relevant policies, the core requirement for document inclusion from the search methods in Phase II was that it contained the word "peat", "peatland", or "wetland" either in English or in the policy's original language. For policies recommended by experts, the source of the recommendation was taken as evidence of the policy's relevance, so some policies were included which did not contain the three core keywords; this exceptionalism introduced investigator bias into the dataset which then has been addressed in Phase III.

For non-English materials, researchers used Google Translate to guide their policy exploration. Language barriers were a main limitation to how deeply or broadly researchers could explore a given country's collection, but this method still functioned to give each country a relatively standardised starting point of peatland policies.

In addition to the keyword search, researchers also combed through government websites and local literature like EU project reports to ensure the collection of other prominent peatland-specific documents missed by the 19 keywords. The government websites of many countries have "site maps" that provide access to topical policy collections unable to be accessed by a general web search; these repositories became a core source for policy documents less prioritised in the algorithms of traditional search engines.

Finally, researchers also relied on the policy framework sections of policies immediately relevant to peatlands (e.g. a peatland strategy) to ensure they were not missing any additional documents pertaining to peatlands. The end of the second phase yielded a total of 405 policies across 27 countries (i.e. UK and all EU member states except Malta) and 22 languages.

[6] https://www.odoo.com/

Phase III (October 2025-Present)

The final phase of policy collection is ongoing. At this point, researchers are focusing more on regional and local policies, requiring the identification and utilisation of siloed local authority and departmental policy repositories. Additionally, the analysts returned to policies from earlier phases to confirm their inclusion or validate their exclusion; after reviewing the expert-recommended policies which were relevant to peatlands but lacked the three core keywords, researchers identified four more keywords indicating a policy's non-explicit relevance to peatlands. This expanded the inclusion criteria as of Phase III, requiring policies to either contain the linguistic equivalents of explicit peatland relevance (i.e. "Peat", "Peatland", "Wetland"), or the linguistic equivalents of non-explicit peatland relevance (i.e. "Organic soil", "Carbon soil", "Water table", "Turf"). This criteria update allowed the inclusion of some formerly rejected policies found via search, and enabled the researchers to address some of the investigator bias introduced by previously holding expert-suggested policies to a separate standard. Policies collected in this phase additionally come from workshops held with experts in different European countries. Future work in AI may enable automatic updates of the policy collection when new documents are released, updated, or expired.

Inclusion Criteria

The policy researchers developed and refined the guidelines in Table 2 throughout the process of dataset creation, finalised as of Phase III.

Table 2. Inclusion and exclusion criteria for the addition of policies to the Policy Collection.

Inclusion Criteria	Exclusion Criteria
Contains linguistic equivalent to "Peat", "Peatland", "Wetland", "Organic soil", "Carbon soil", "Water table", or "Turf".	Does not contain any linguistic equivalent to "Peat", "Peatland", "Wetland", "Organic soil", "Carbon soil", "Water table", or "Turf".
Published by a democratically elected governing body or its department.	Published by an NGO or a third party with no participation of democratically elected governing body or its department.
Published as a written policy document.	Published as infographic, recording, etc.
Policy documents relating to decision-making and strategic planning.	Legal documents relating to the enforcement of laws and regulations, or administrative documents like agendas, meeting minutes, reports, etc.
Published in one of the official languages in the EU.	Published in a language not recognised as an official language in the EU.

Edge Cases

In order to demonstrate the inclusion criteria in action, Table 3 provides examples of 4 policies on the border of being included in the Policy Collection as well as the logic behind the final decision to include or exclude them from the dataset.

Information Recorded

For each policy, the following information is recorded for future display in the Policy Collection page and Policy Profile pages: Name of policy in English, Name of policy in original language, Administrative level, Beginning and end year of policy duration, Country (if national or subnational), Region (if subnational), Language of policy text, Keywords, Publisher, Link to the document PDF, Link to the policy page (if exists), Excerpt of text that proves peatland relevance (in English and original language), Overview of policy from text or policy page (in English and original language), and Table of Contents (in English and original language).

Categories

The policy categories detailed in Figure 2 were developed prior to Phase I by evaluating a selection of peatland-affecting policies for key themes. This analysis yielded eight core categories of policies relevant to peatlands. The policies are labelled with a single primary category, then all other relevant categories. These categories are assigned by single coders for the purpose of providing general context for the policy topic, not for the purpose of analysis.

Table 3. Policies on the edge of being included or excluded from the Policy Collection.

Country - Policy Title	Reason for Uncertainty	Included (Y/N)	Decision Explanation
Malta – Common Agricultural Policy Strategic Plan – Malta, 2023–2027	Malta does not have any peatlands.	Y	The keyword 'peatland' appears twenty-two times. Even though there are no peatlands to be affected in Malta, the policy meets the inclusion criteria, so it is included.
France – National Low Carbon Strategy	The document does not mention peat, peatlands, or wetlands, so it does not appear to directly affect peatlands.	Y	Although the policy does not mention peat, peatlands, or wetlands, it mentions 'carbon soils'. The policy affects French peatlands indirectly, so is included.
Germany – Lower Saxony Act to Promote Climate Protection and Mitigate the Consequences of Climate Change (NKlimaG)	The document governs peat soils but is an 'Act' instead of a policy document related to decision-making or planning.	N	This document was excluded because it is piece of formal legislation (law) and not a policy document related to decision-making or planning.
Ireland – Chapter 6: Soil Environmental Protection Agency's 2024 State of the Environment Report	This document contains references to peatland, peat, and wetlands, but it is a chapter of a report.	N	Although the document met the terminological inclusion criteria, it was excluded as it is an administrative report instead of a policy document related to decision-making or planning.

Characteristics

As of this paper's publication, the current collection of European peatland policies consists of 571 policy documents across 28 countries and 22 languages. The countries include all EU member states and the United Kingdom. The languages include all official EU languages except Irish and Maltese. This section explores the composition of the Policy Collection by language, by country for national and sub-national policies, by governance level, and by primary category.

Language Composition

Figure 3a shows the makeup of the Policy Collection by document language. As researchers were directed to default to collecting the official English translations of policies when available, English is overrepresented in the dataset. Following English, German is the most common document language, correlating with the large number of policies from Germany in the dataset. The French and Dutch language counts include Belgian French and Flemish language policies, respectively.

Country Composition

Figure 3b provides a visual illustration of the composition of the Policy Collection by country. German policies are the best represented in the dataset, followed by Ireland, Italy, Slovakia, and the United Kingdom. Malta is the least represented country in the collection as it does not have any peatlands, but it does still contain some peatland policies; the logic for their inclusion can be found in Table 3.

Governance Level and Category Composition

The composition of the Policy Collection by governance level and primary category is illustrated in Figure 4. The categories assigned to the policies are single-coder, so they should be taken as a general indication of policy topic instead as a reliable basis for analysis. National policies represent the vast majority of the collection as regional and local policies were deprioritised in the initial phases of policy acquisition. In terms of thematic distribution, researchers set "Biodiversity" as the primary category most frequently, followed by "Land Use and Agriculture", "Climate Action", and "Environmental Quality". The structure of the acquisition method and researcher judgments did not yield many peatland policies with primary topics of "Research and Applied Sciences" and "Community and Culture".

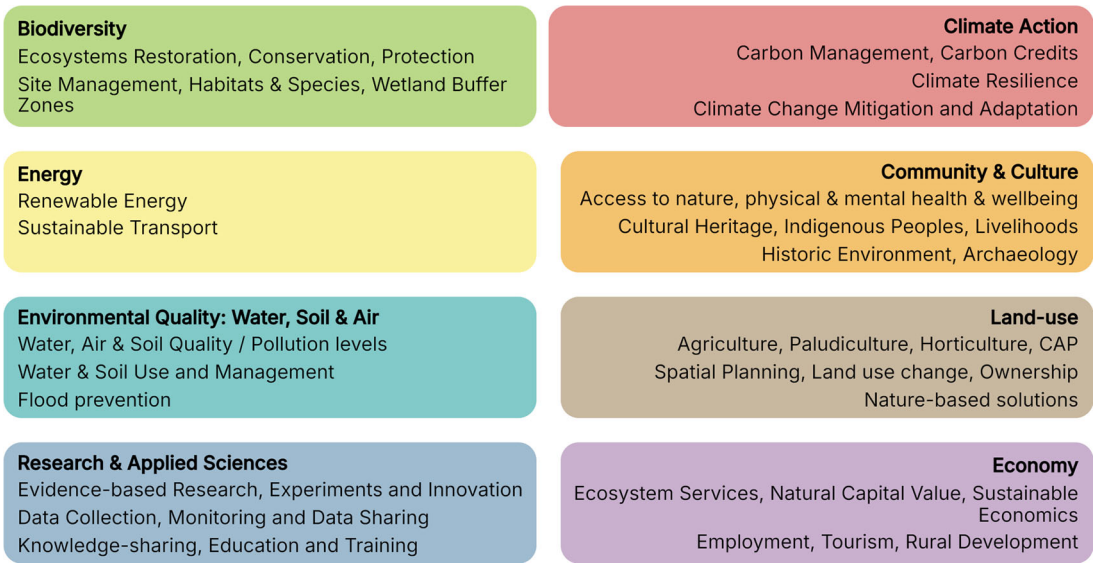


Figure 2. The names and definitions of the eight policy categories used to classify the policies of the European Peatland Policy Dataset.

Comparative Coverage Analysis with Overton

To validate the scope and necessity of the peatland policy dataset, researchers conducted a quantitative cross-database analysis between the most recent version of the portal's Policy Collection (571 total policies, of which 541 were at national and subnational levels), and results retrieved from Overton [4] with the search conditions below (174,229 total documents, 117,813 national and subnational). Overton's pool of results was much larger than that of the portal's due to its inclusion of additional grey literature alongside the actionable plans and strategies in the scope of the portal, so researchers focused their analysis on the presence of the portal's policies in Overton's dataset.

The search conditions to retrieve results from Overton consisted of:

- Search String: peatland OR peat OR wetland OR carbon soil OR organic soil OR water table OR turf
- Source Sector: Public sector

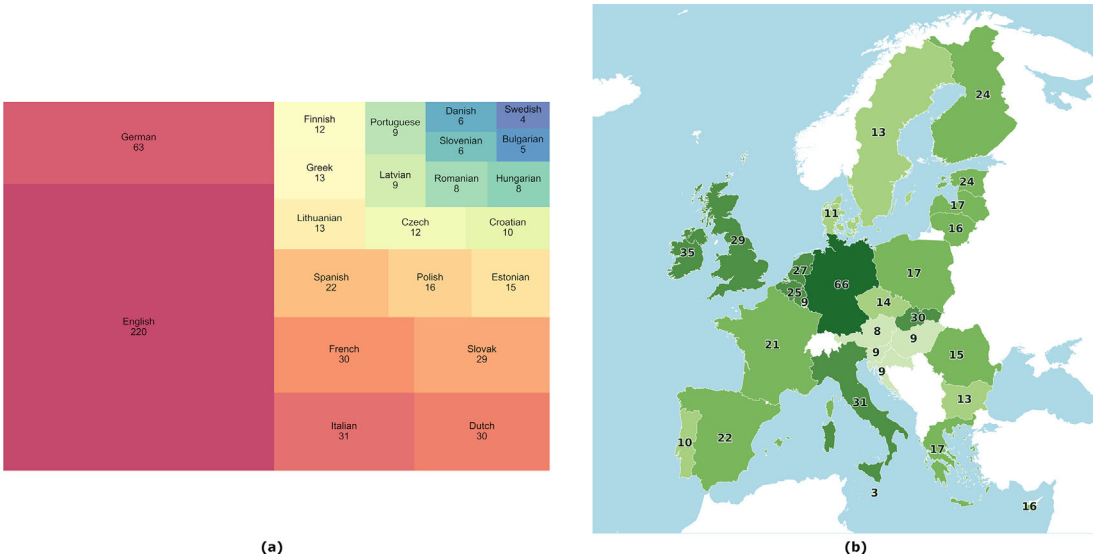


Figure 3. Graphs representing the composition of the Policy Collection by language and country. (a) A treemap showing the number of policies of each language in the Policy Collection. (b) A choropleth map showing the number of national and subnational policies from each country in the Policy Collection.

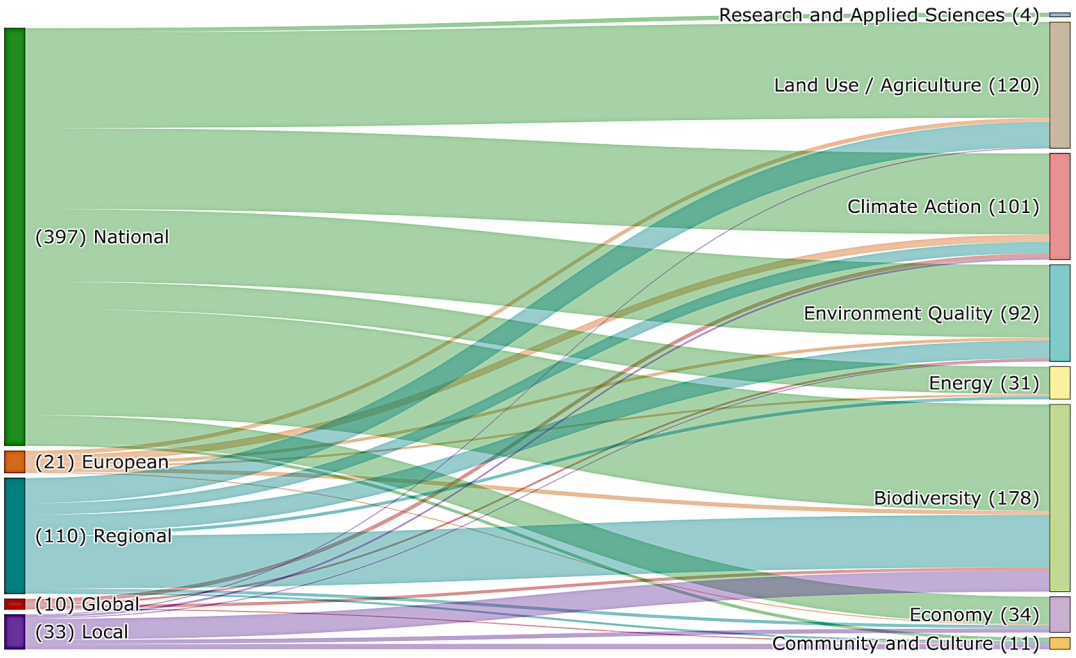


Figure 4. A Sankey diagram demonstrating the makeup of the Policy Collection across governance levels and primary categories.

- Source Organisation Type: Government, Legislative Body, Judicial Body
- Years: 1959 to 2026 (the portal's policies begin at 1960)
- Source Country: Every EU country and UK
- Document Type: Publication
- Search Date: 21st July 2026

The search string was derived from the inclusion criteria's seven keywords, as Overton allows users to search in the body text of documents as well as the title. The search filters of Sector, Organisation Type, Years, Country, and Document Type were applied in an attempt to limit the results to the same kinds of documents as found in the portal.

The coverage analysis was conducted by calculating the Levenshtein distance (i.e. text string overlap) between titles from both policy datasets in order to identify shared entries. As Overton does not include a field to select global policies, researchers limited their analysis to the national and subnational policies contained in each dataset.

A two-pass matching system was used to account for language difference. The system first tried to match policies using their English titles; if it failed, it attempted the same with titles in their original language. Finally, a hand-tuned 95% similarity threshold was set to minimise false positives after several diagnostic trial runs.

Quantitative Overlap

A summary of the overlap between the Overton dataset and the portal's dataset at national and subnational levels can be found in Table 4, where the reported match percentage indicates the proportion of the portal's policies found in Overton's collection. Overall, the Overton search retrieved only 74 (13.68%) of the portal's 541 national and subnational policies. The highest alignment rate by far was present within the United Kingdom (48.28%), possibly due to overrepresentation of English language policies or language bias in the collection protocols of both the portal and Overton. After the United Kingdom, Bulgaria (38.46%), Ireland (34.29%), and Austria (33.33%) had the highest overlap, while Cyprus, Greece, Hungary, Lithuania, Luxembourg, Malta, and Slovenia had no matches at all.

Qualitative Analysis: Luxembourg

In order to explore the retrieved policies from each dataset, researchers performed a qualitative analysis of Luxembourg's policies. To achieve a numerically comparable pool of policies between Overton and the portal, an Overton search was performed as previously described except that the search string now consisted of the 3 primary keywords instead of the 7 used previously. With this change, Overton returned 12 documents for Luxembourg, matching again 0 of the portal's 9 documents; the document names for each dataset are listed in Table 5. As expected, the results from Overton contain documents outside the scope of the portal such as reports, agendas, and debate proceedings. Additionally, the presence of non-descriptive "Bill", "Consolidated file", and "Document name" results reveal indexing limitations of automated web search protocols. Diagnostic testing showed that relaxing the similarity threshold to 90% falsely paired two sets of unrelated documents, reinforcing the validity of the 95% threshold to avoid false positives.

Overton's omission of the vast majority of the portal's policies, likely due to algorithmic blind spots, suggests additional value in the portal's use of expert input through workshops and partner contributions to construct the dataset. Further, while Overton's sheer volume means that it may contain relevant documents currently absent from our collection, its

Table 4. Matches between policy document and grey literature titles from Overton and the portal, limited to national and sub-national levels. Match percentages over 30% are in bold.

Country	Portal Total	Overton Total	Match	Portal Match (%)
Austria	9	2900	3	**33.33**
Belgium	25	3029	1	4
Bulgaria	13	1349	5	**38.46**
Croatia	9	592	1	11.11
Cyprus	16	742	0	0
Czech Republic	14	470	1	7.14
Denmark	11	1560	1	9.09
Estonia	24	1312	4	16.67
Finland	24	2729	2	8.33
France	21	3263	1	4.76
Germany	66	3610	9	13.64
Greece	17	526	0	0
Hungary	9	321	0	0
Ireland	35	12193	12	**34.29**
Italy	31	1645	1	3.23
Latvia	17	1243	4	23.53
Lithuania	16	123	0	0
Luxembourg	9	340	0	0
Malta	3	212	0	0
Netherlands	27	3773	2	7.41
Poland	17	177	1	5.88
Portugal	10	1065	1	10
Romania	15	755	2	13.33
Slovakia	30	657	3	10
Slovenia	9	798	0	0
Spain	22	5162	2	9.09
Sweden	13	7505	4	**30.77**
United Kingdom	29	59762	14	**48.28**
Total	541	117813	74	13.68

Table 5. List of the retrieved English titles from Overton and the Policy Collection, limited to Luxembourg.

Overton	Portal
2020 activity report of the National Museum of History and Art - government.lu	CAP - National Strategic Plan of the Grand-Duchy of Luxembourg
Agenda	Draft strategy and action plan for adaptation to the effects of in Luxembourg (2025–2035)
Bill	Information document relating to the 3rd National Plan concerning the Protection of Nature
Bill	Integrated National Plan on energy and climate change in Luxembourg for the period 2021–2030
Consolidated file	Luxembourg Climate Bonus Forest Moor and Meadow subsidy
DEBATE ON THE POLICY OF HUNTING	National Forestry Accounting Plan Luxembourg
Document name	National Nature Conservation Plan - 3rd Plan to 2030
Document name	National Plan for the Protection of Nature
Documentation Bulletin No. 15/1965	Third Management Plan For The Luxembourg Portions Of The International River Basin Units Rhine And Meuse (2021–2027)
Official Journal of the Grand Duchy of Luxembourg	
Renaturation of waterways - Restoration of wetland habitats - government.lu	
Special Bulletin November 1980	

inclusion of additional grey literature and nondescriptive metadata highlights a trade-off in utility; for users looking to identify policy actions and instruments without wading through excess administrative documents, the portal's collection is an efficient, publicly-available alternative.

Policy Exploration Tools

The goal of the portal was not just to compile these scattered policies, but also to provide interfaces for policy researchers to explore the dataset. This resulted in the project producing three main avenues of interaction with the collection: first, a simple tabular interface for viewing and filtering the policies, as well as more-detailed policy profiles for each document; next, a map interface which shows the peatlands and peat soils of the EU and, with the click of a mouse, reveals all the policies that affect a single point; and finally, a RAG-based question-answering tool built on the full texts of all the policy documents, allowing users to ask questions and receive responses alongside the source documents of each answer.

Policy Collection

The Policy Collection provides a structured, tabular interface for exploring the European Peatland Policy Dataset (Figure 5a). The collection is filterable by country, level, and category, as well as searchable. Each policy additionally has its own profile page with more detailed overviews of the policies including the original PDF document and contribution mappings to both the EU Green Deal's goals and the UN's Sustainable Development Goals (Figure 5b).

The design of this interface is typical for policy research: lists of policies with a few key details included like Title and Location; filters for location, level, and category; a search function matching titles and overview text to find relevant documents. The portal presents the policy dataset in this format with the goal of streamlining the most common policy discovery tasks. Though created with policy researchers in mind, the page is designed to help any user explore the catalogue.

Peatland Policy Map

While a catalogue structure is the traditional way to explore a policy dataset, the portal's Peatland Policy Map provides a spatial view of the peat soils and peatland policy landscape of Europe. This tool allows any user to explore which policies affect which peatlands, and aims to provide policy researchers with unique insights into the peatlands requiring policy attention. Figure 6 provides an example of the interface when a point on the map has been selected.

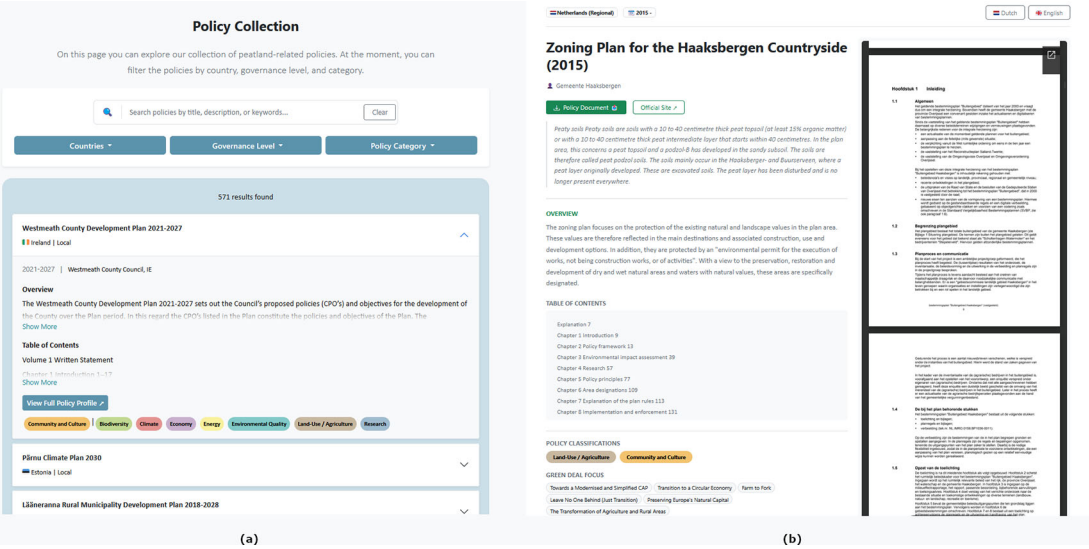


Figure 5. The portal's Policy Collection page. (a) The page's tabular list view of all policies, searchable and filterable by country, governance level, and category. (b) An example of the profile page view for each policy, containing any additional information and resources recorded in the dataset.

The Peatland Polygons

There is no official, complete map of the peatlands of Europe, largely due to the persistent lack of a standardised definition of what constitutes a peatland,[27] as well as the lack of a targeted effort throughout Europe to record peatlands in a unified, systematic way. Tanneberger et al. (2017) produced a peatland map of Europe, estimating peat soil coverage from various types of national datasets and proxy variables,[1] and more recently Tegetmeyer et al. (2025) released the European Wetland Map (EWM), also pieced together using national datasets with differing types of data sources.[28] Our peatland map of Europe also follows in this vein, using national datasets of peatlands if they exist,[29] otherwise using sources like general soil maps or topographic databases. When the EWM was released, the portal updated some of its national peatland maps to use the EWM data, filtering it to only include wetlands specifically designated as "peatland." A full list of the source datasets and their transformations can be found on the portal [7].

Initially, the map included polygons for countries from the CORINE Land Cover (CLC) [8] dataset of the EU's Copernicus Programme, filtered to only include class 412 "Peatbogs." We removed this polygon, however, due to its low resolution and its exclusion of converted peatland soils to other land cover classes (e.g. forest, or grassland).

When a user selects a peatland on the map, a small pop-up displays information about the chosen polygon. Due to a lack of cohesive structures across peatland datasets, the information fields provided in the pop-up differ across data sources and even between countries of the EWM. Despite this, most polygons will at least include the type of peatland present.

The display of the peatland polygons in their original, full-resolution form caused latency issues in the map due the sheer amount of data being loaded by the server. In order to address this, scale dependencies were introduced so that very small polygons were hidden and the rest of the polygons were simplified when zoomed out on the map. These changes, and the lack of in-house validation of the polygons through ground-truthing, mean that the peatlands displayed on the map should not be taken as authoritative peatland maps, and instead as approximate land-cover indicators for context in understanding the peatland policy landscape.

The Policy Polygons

While the policy polygons are invisible to the user during navigation, they serve as the underlying spatial triggers for the map interface. By georeferencing national and sub-national policies to the territorial boundaries of their respective

[7] https://peatlandpolicyportal.eu/map-info
[8] https://land.copernicus.eu/en/products/corine-land-cover



Figure 6. The Peatland Policy Map enables hyper-local spatial analysis of the policy landscape alongside visualisations of recorded peatlands.

regions, the portal defines each document's specific geographic extent. When a user selects a coordinate on the map, a sidebar populates with all applicable policies, organised in ascending order of governance level: local, regional, national, EU, and finally, global policies. This hierarchical display, combined with integrated filtering tools, is designed to allow researchers to isolate specific administrative levels or thematic categories to analyse the multi-layered regulatory environment of a single point. This method of exploring the policy landscape spatially is additionally intended to be used to investigate the conflicts noted in Chen et al. (2023).[4] Further, as peatlands are ecosystems which do not respect administrative boundaries, this spatial visualisation is also intended to identify gaps where policies are only in force for one small region of a peatland ecosystem. Ultimately, the goal of this tool is to help policy researchers to understand the peatland policy landscape in a given area and identify policy conflicts and gaps to be addressed.

There are some limitations to assigning policies spatial boundaries by administrative extent, as some policies target only specific sites (e.g. Natura 2000 sites) or specific land types (e.g. organic soils) within a region. For example, the Dutch Natura 2000 Management Plan for Witte Veen (*Natura 2000 Beheerplan Witte Veen*) only affects the specific Witte Veen site, but the policy is assigned the same administrative boundaries as the publisher, the Province of Overijssel; when a user clicks anywhere within the Province of Overijssel, even outside the boundaries of Witte Veen, this policy will be retrieved and listed as affecting that location. The flattening of policy sub-scoping through use of general administrative boundaries to define spatial effect does introduce false positives to the retrieved policy lists, so users are alerted to this limitation in the Policy Map's explanation page [7].

RAG Question-Answering

The final method of interaction with the policy dataset is in the form of an AI Question-Answering tool. Any user can interact with the tool by submitting natural-language questions about peatland restoration policies in any major European language, without requiring prior knowledge of the original policy language or jurisdiction (Figure 7). Upon receiving a query, the system retrieves and analyses relevant evidence from the Policy Collection's PDFs, generating a response that is explicitly grounded in the source documents.[30] This approach addresses the phenomenon of hallucinations in large language models (LLMs), where models generate plausible but incorrect content.[31] RAG's use of the retrieval step forces the LLM to constrain its generated response to the evidence in its library, increasing the reliability of the model.[32] The generated answer is accompanied by transparent citations and a ranked list of the most relevant policy documents, enabling users to verify the provenance of each claim and explore the underlying political or strategic texts directly.

Beyond basic question answering, the system provides fine-grained control over the retrieval scope through metadata-based filtering. Users may refine the policy dataset used for answering by specifying constraints such as governance level (e.g. EU-wide directives, national strategies, regional plans), policy language, or country of origin. These filters are applied consistently across both dense semantic retrieval and lexical search components, ensuring that only contextually and jurisdictionally appropriate documents contribute to the final response. This capability is particularly important in peatland governance, where similar policy goals may be articulated differently across countries, languages, and regulatory levels.

Together, multilingual querying, hybrid evidence retrieval, explicit citation, and user-controlled filtering allow the AI tool to support trustworthy, transparent, and context-aware policy discovery. The system is designed not only to answer questions, but also to assist researchers, policymakers, and practitioners in navigating complex, multilingual policy landscapes while maintaining traceability to authoritative sources.

RAG System Architecture

The specifications for the RAG tool are as follows:

Document ingestion and chunking. Policy PDFs are parsed using PyPDF and segmented into overlapping chunks of 500 characters with 50-character overlap using LangChain's RecursiveCharacterTextSplitter. Each chunk is stored in a Qdrant vector database with metadata fields: filename, country, language, governance level, category, and record ID. All metadata values are normalised to lowercase at ingestion and query time to ensure consistent filtering.

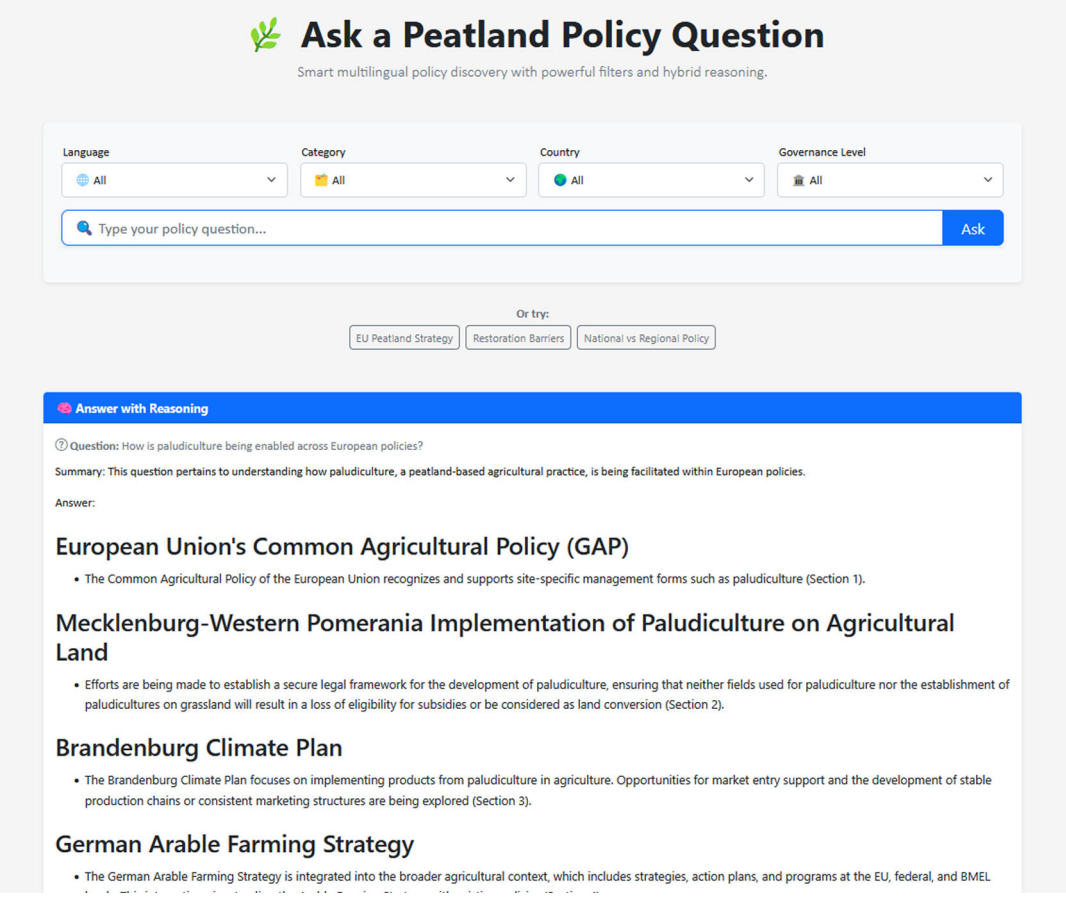


Figure 7. The RAG-based question-answering tool allows users to query the entire policy dataset and receive both a response and sources in seconds.

Embedding model. Dense embeddings are computed using paraphrase-multilingual-mpnet-base-v2 (768 dimensions),[33] a sentence-transformer model supporting 50+ languages. This model enables queries in one EU language to retrieve relevant chunks from documents written in another - a core requirement given the 22-language corpus.

Hybrid retrieval. The system combines dense semantic retrieval with BM25 lexical scoring. The hybrid score is:

$$s(d|q) = \alpha \cdot \cos(\varphi(q), \varphi(d)) + \beta \cdot \mathrm{BM25}(q, d)$$

where φ denotes the sentence embedding function. The top k = 10 candidates are retrieved above a minimum score threshold of τ = 0.3. Results are deduplicated by source document before generation, ensuring diverse evidence from multiple policies.

Metadata filtering. Users may apply filters on language, country, category, or governance level. Filters are applied as must-match conditions in Qdrant prior to vector search, ensuring retrieved chunks are both semantically relevant and jurisdictionally appropriate.

Multilingual handling. The query language is detected automatically using langdetect (deterministic seed = 0), covering all 24 official EU languages. The LLM is instructed to respond in the detected query language regardless of the source document language, enabling cross-lingual policy synthesis. Pre-defined not-found messages exist for all 24 EU languages.

Generation parameters. Answers are generated by a locally hosted Mistral 7B model served via Ollama (temperature: 0.2; context window: 4,096 tokens; maximum output: 600 tokens). The prompt contains: (1) a system instruction to respond only from provided context and in the detected query language; (2) the top retrieved chunks labelled by source and author, truncated to 6,000 characters; and (3) the user query.

Citation grounding. Each response is accompanied by a ranked list of source documents ordered by hybrid retrieval score, with metadata including filename, author, country, language, and score. When no chunks exceed the relevance threshold, the system returns a pre-defined not-found message in the query language rather than generating an unsupported answer.

Quantitative Evaluation: RAGAS

To provide an empirical assessment of the RAG tool, researchers constructed a held-out evaluation dataset of 58 question–answer pairs generated automatically from 12 peatland-specific policy documents in the indexed corpus using the RAGAS TestsetGenerator framework (ragas v0.1.21).[34] The dataset spans nine countries (Belgium, Bulgaria, Germany, Greece, Ireland, Latvia, Romania, Netherlands, and the United Kingdom), four governance levels (supranational, national, regional, and local), and three document languages (English, German, and Dutch). Questions cover four types: Targets/Numeric (n = 24), Dates (n = 14), Purpose/Definition (n = 13), and Location (n = 7). Questions and ground truth answers were generated by Mistral 7B from real policy document chunks, grounded in actual indexed content but model-generated rather than manually annotated by domain experts, which is standard practice for automated RAG evaluation at scale.[34] The dataset and generation scripts are archived on Zenodo.[35]

Three RAGAS metrics are reported. Context precision measures whether retrieved document chunks are relevant to the question, operationalising retrieval quality. Answer relevancy measures whether the generated response addresses the question. Citation precision is a threshold-based measure of source quality computed without an LLM. Classical IR metrics (recall, precision) are not applicable at scale here, as per-query relevance labels do not exist for this 421-document, 22-language corpus; context precision serves as the accepted proxy.[34,36]

As shown in Table 6, the system achieves a mean context precision of 0.895, indicating that in approximately 9 out of 10 retrieval operations, the returned document chunks are genuinely relevant to the question. Precision is highest for date queries (0.978) and numeric target queries (0.940), consistent with the hybrid BM25 and dense retrieval architecture's strength in lexically precise questions. Location queries show lower context precision (0.657), reflecting the limitation of chunk-level retrieval without direct integration of the portal's spatial metadata — a direction for future development. Mean answer relevance is 0.638, with purpose and definition queries achieving the highest score (0.828).

Table 6. RAGAS evaluation results across 58 held-out questions.

Question Type	N	Answer Relevance	Context Precision	Citation Precision
Dates	14	0.642	0.978	1.000
Targets/Numeric	24	0.614	0.940	1.000
Purpose/Definition	13	0.828	0.851	1.000
Location	7	0.406	0.657	1.000
Overall	**58**	**0.638**	**0.895**	**1.000**

Quantitative Evaluation: Baseline Comparison with non-RAG LLM

To provide a direct baseline comparison, the same 58 questions were posed to Mistral 7B in zero-shot mode (i.e. without any retrieved document chunks) and answer relevancy was computed using the same RAGAS framework. Table 7 reports results for both systems, and Table 8 provides example responses from both as well. Zero-shot Mistral achieves a marginally higher mean answer relevancy (0.684 vs 0.638), however this difference requires careful interpretation: answer relevancy measures only whether a response addresses the question, not whether it is factually correct or traceable to an authoritative source.

Critically, zero-shot Mistral produced confident responses to all 58 questions, including highly specific queries about regional policy provisions — for example, the exact habitats listed in a 2009 Romanian peatland restoration annex or the stakeholder engagement process of a specific national wetland plan — that cannot be reliably answered from parametric knowledge alone. None of the 58 zero-shot responses acknowledged uncertainty or referred the user to a source document. This is the hallucination risk that RAG is designed to address.[37,38]

By contrast, the RAG system returns a pre-defined not-found message when no relevant chunks exceed the retrieval threshold, prioritising factual honesty over fluency. Every response produced by the RAG tool is accompanied by a ranked list of source documents, enabling users to verify each claim against the original policy text. In the context of evidence-based policymaking, where a plausible but incorrect answer about a restoration target or a legal deadline could have real consequences, this traceability is the system's primary value, not its answer fluency score.

Hallucination Control and Limitations

A formal hallucination audit by independent raters was not conducted and is acknowledged as a limitation. The portal's citation mechanism, whereby every generated response is accompanied by a ranked list of source documents, provides a user-verifiable proxy for hallucination control. A systematic expert audit is planned as future work.

Additional Tools

The aforementioned tools are the core contributions of the portal to the area of peatland policy analysis, but the portal has additional tools offering reference for policy researchers and other specialisations.

Project Catalogue

The Project Catalogue showcases the portal's collection of European peatland projects. While the vision is for the catalogue to encompass projects down to the local level, the first stage of building the collection has focused primarily on peatland-affecting EU LIFE projects from the last 15 years.

The portal defines its peatland projects of interest to be those which are concerned with practical restoration, conservation, and monitoring of peatlands. This includes projects ranging from drain-blocking, a common method of rewetting peatlands, to paludiculture, a new agricultural practice of wetland farming on rewetted peatlands. Most peatland restoration projects receive financial support from governmental institutions, often through a funding mechanism from a policy or national implementation of an EU directive. The EU LIFE program that funds the majority of projects in the catalogue is managed by the EU's European Climate, Infrastructure and Environment Executive Agency (CINEA) [9] which manages the grants that fund projects to support the goals of the EU Green Deal. This project collection is designed

[9] https://cinea.ec.europa.eu/

Table 7. **RAG tool versus zero-shot Mistral 7B across 58 questions.**

Metric	RAG Tool	Zero-Shot Mistral 7B
Answer Relevancy	0.638	0.684
Context Precision	0.895	N/A — no retrieval
Citation Precision	1.000	N/A — no sources
Responses citing source documents	58/58	0/58
Refusals when evidence absent	**Yes**	**Never**

Table 8. **Representative examples comparing RAG tool and zero-shot Mistral 7B responses.**

Question Type	Question	RAG Tool Response	Zero-Shot Mistral Response
Dates	What part does the NATURA 2000 network of SPAs play in safeguarding endangered & migratory birds amidst threats?	Retrieved from: vassilevetal2013Wetlands-planEngl.pdf (Bulgaria). The NATURA 2000 network of SPAs plays a significant role in safeguarding endangered and migratory birds by providing habitat protection through a coherent network of Special Protection Areas, as established under the Habitats Directive. [Source cited with retrieval score]	The NATURA 2000 network plays a crucial role in safeguarding endangered and migratory bird species across Europe. These areas are designated under the EU Birds Directive (79/409/EEC)... [No source document cited. Parametric knowledge only. Answer relevancy: 0.762]
Targets/ Numeric	What are the natural habitats of communal interest in Romania regulated by the Habitats Directive belonging to the category of bogs and peatlands?	Retrieved from: Annex 10_Guide for the restoration of degraded peatlands from Romania.pdf. The habitats are those listed in Annex 1 of the Habitats Directive, as transposed in Romania through OUG No. 57/2007 (approved by Law No. 49/2011). [Source cited with retrieval score]	The Habitats Directive regulates several natural habitats of communal interest in Romania, including Ombrotrophic raised bogs (H50), Transition mires and quaking bogs (H7140), and Calcareous fens... [Habitat codes and details generated from parametric knowledge — not verified against Romanian legislation. No source cited. Answer relevancy: 0.709]
Purpose/ Definition	What is the typical behavior of a Pygmy Cormorant during winter, given its status as a protected species under Ramsar criteria?	The requested information is not found in the provided documents. [System correctly identifies the question is outside the scope of the indexed policy corpus and returns a pre-defined not-found message in English.]	The Pygmy Cormorant (Microcarbo pygmeus) is a protected species under the Ramsar Convention. During winter, these birds are typically found in coastal areas and large lakes with open water, as they prefer shallow waters for feeding... [Fluent but entirely from parametric knowledge. Unverifiable in policy context. Answer relevancy: 0.763]

Table 8. *Continued*

Question Type	Question	RAG Tool Response	Zero-Shot Mistral Response
Location	What are the potential pros and cons for long-term eco-friendly management of moorlands, taking into account atmospheric pollution?	Retrieved from: national_wetland_policies_-_greece.pdf. The pros include preservation of natural moors from degrading due to atmospheric pollution (acid rain, airborne nutrients). The cons include water scarcity and rising temperatures. [Source cited with retrieval score]	Pros: Carbon sequestration, biodiversity, water regulation, cultural heritage. Cons: Economic loss for landowners, difficult terrain, climate change vulnerability, high restoration costs... [Generic list from parametric knowledge. No policy document cited. Answer relevancy: 0.727]

to provide policymakers with examples of practical solutions and actions to include in their policy proposals for supporting peatlands.

The catalogue interface itself, as seen in Figure 8a, is a simple tabular structure like the Policy Catalogue. Each project entry contains the project title and official project code in a simplified information view, with dates of operation, partner countries, project site names, and a link to the project page available upon selection of the project.

Project Map

The Project Map provides a spatial representation of the peatland project landscape in Europe. This view of the project collection allows policymakers to see where projects are occurring, identify gaps in project coverage, and examine the diversity of projects active within their own administrative regions.

This map, like the Peatland Policy Map, uses OpenStreetMaps as its basemap, then displays all project sites from the catalogue distributed geographically. Projects may have more than one site in more than one country, so the sites for each project all appear on the map first as grey location icons, accompanied by a sidebar which lists all projects.

Selecting a project site on the map triggers a dynamic visual query: the chosen icon transitions to red, while all associated sites within that project are highlighted in pink to indicate their connection. Simultaneously, the sidebar populates with detailed project metadata, including operation dates and partner countries. This interface is intended to allow users to evaluate specific site data while maintaining a visual overview of the project's broader geographic footprint, facilitating a better understanding of multi-site restoration initiatives. This view is provided in Figure 8b.

Area Information Pages

The portal provides area information pages in the form of country factsheets containing information about the peatlands, projects, and most important peatland policies in that area. These pages are designed to allow users from any domain to get an efficient, informed summary of the peatland and policy landscape for a given area.

The peatland facts provided include an area's total peatland coverage, and of that, its natural, cutover, forested, or agricultural peatland coverage. As definitions of peatlands vary however, even within countries[27] and between organisations reporting different values of each type of peatland coverage, the portal is scrupulous in its fact management. Instead of arbitrarily choosing one correct value of peatland coverage, the area information pages instead provide any reliable figures, if they exist. Alongside the value and unit for a certain statistic, the page also provides a link to the source and the date of the fact's publication (Figure 9a). This transparency allows users to see the range of values, definitions, and methods used in quantifying peatland coverage, and to decide for themselves which figures to use in their own work.

The key policies in a factsheet are selected subjectively by policy researchers in the factsheet's area, and are intended to be a simple overview of the most influential or notable policies in an area. Those users looking for the full list of an area's policies contained on the portal can filter the Policy Collection page to achieve that goal.

Best Practise Reference

The other general reference tool provided by the portal is the Best Practise Reference page (Figure 9b). The page's topics are separated into "General Policy", aimed at policymakers looking for best practice in policy development, "Restoration", for users interested in best practice for practical peatland management and restoration tasks, and

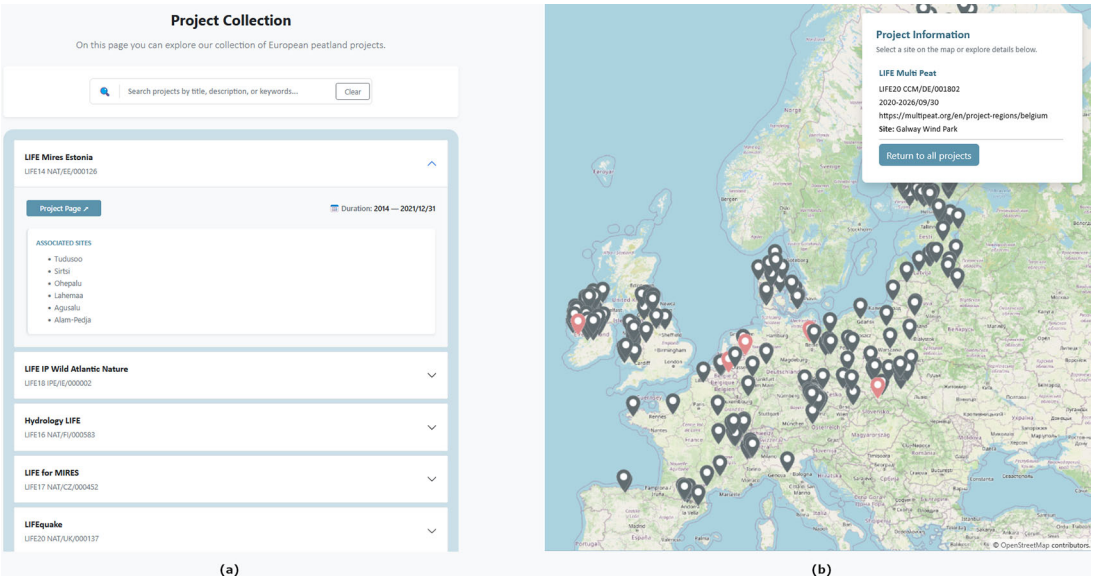


Figure 8. The Project Catalogue page. (a) The tabular list view of all projects, searchable and filterable by country. (b) The map display of projects across their various sites.

"Standards", to encourage researchers and experts in the peatland domain to work from centralised terminologies and measurement methodologies. The Best Practise Reference page is intended to grow as more users share their recommendations, with special attention given to the policy reports from EU projects.

Restoration Calculators

The Restoration Calculators evaluate peatland restoration scenarios for efficacy and financial feasibility. These tools are mostly aimed at land managers and project managers, but they additionally have the potential to provide policy researchers with an understanding of the financial scope of restoration projects funded by EU policies. The tools were developed prior to the Multi Peat project, however their initial format made them inconvenient and uninviting for non-experts. The portal adapted these tools to an online format to make them accessible both in use and in comprehension.

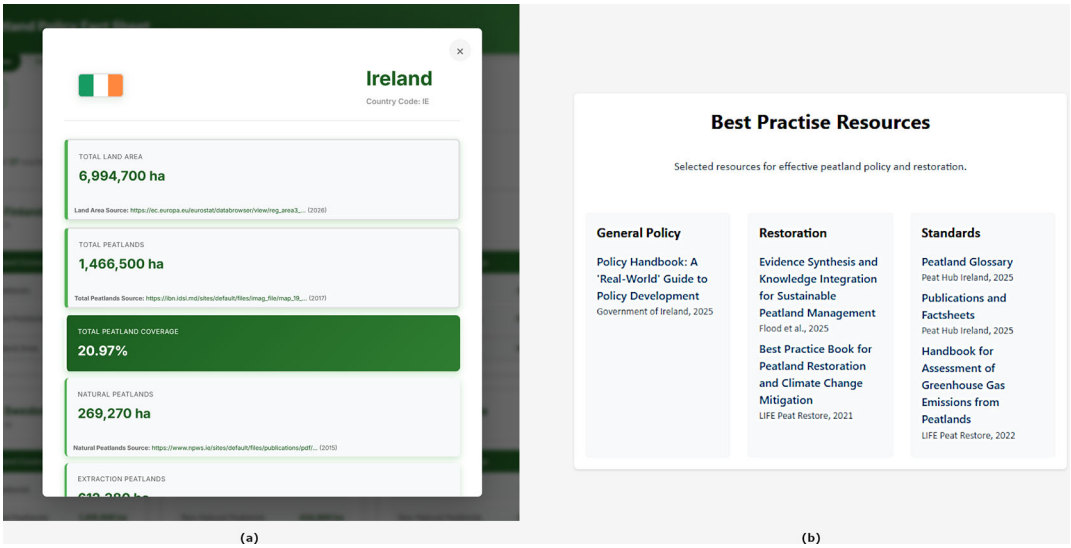


Figure 9. The portal's general reference pages. (a) Ireland's Factsheet. (b) Best Practise Reference page.

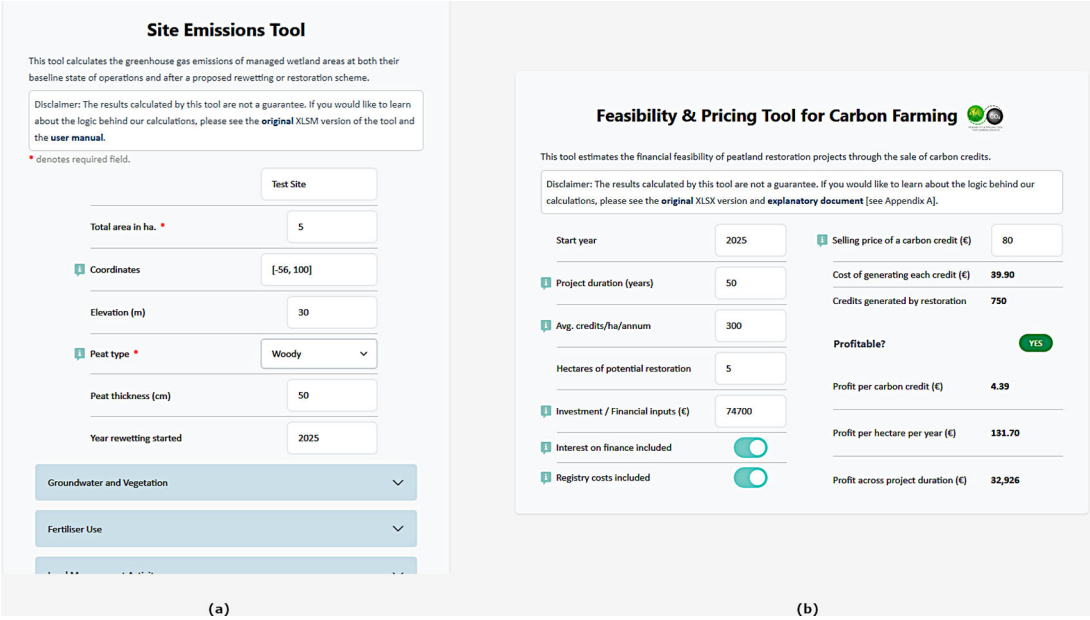


Figure 10. Restoration calculators available on the portal. (a) The Site Emissions Tool,[39] adapted to the web from an XLSM file. (b) The Financial Feasibility and Pricing Tool,[40] adapted to the web from an XLSX file.

Site Emissions Tool

The Site Emissions Tool, developed in the EU Interreg Carbon Connects project, uses the Greenhouse Gas Emission Site Types (GEST) methodology[10] to calculate the estimated GHG emissions of a managed agricultural peatland both before and after a water level, crop, or other management change.[39] This helps land managers predict the tonnes of GHG per hectare per year that would be saved by a proposed change in the peatland's management, which is equivalent to the carbon credits that would be generated by the restoration. While the tool currently provides an adequate estimate, the GEST method is limited to measurements from sites in continental Europe; as such users are made aware that the application of this estimator outside this region may under or overestimate emissions. In the future, the tool will be updated to allow users to select their region for a more accurate estimate. Figure 10a shows the tool and the information required for emissions estimation, including peat type and thickness, groundwater level and vegetation to determine the site type, management information including grazing animals and growing crops.

Financial Feasibility and Pricing Tool

The Financial Feasibility and Pricing Tool, or the Carbon Farming tool, was developed in the EU Interreg Care-Peat project to provide an estimate of how financially feasible a peatland restoration project would be, based on the number of carbon credits the site would be able to generate and sell per hectare each year.[40] The tool has an assumptions tab which can be updated to tune values such as the rate of inflation, but the main part of the tool can be seen in Figure 10b. The

Table 9. How the portal's tools assist in addressing policy challenges.

Policy challenge	Portal tool	Desired policy outcome
Fragmented multi-level policy landscape	Policy catalogue	Rapid identification of all applicable peatland policies across governance levels
Regulatory overlaps and gaps	Peatland policy map	Detection of policy conflicts, gaps, and cross-border inconsistencies
Limited accessibility of policy evidence	RAG-based QA tool	Transparent, citation-grounded policy synthesis and exploration
Difficulty transferring best practices	Project & best practice catalogues	Knowledge transfer across regions and policy contexts

essential data needed for calculation is the project duration, the number of carbon credits generated per hectare per year which can be obtained from the Site Emissions tool, hectares of restoration, financial investment, and the selling price of a carbon credit in the region. Policymakers can explore this tool to develop an understanding of how carbon credits, a key mechanism of blended finance, can be used to fund peatland restoration outside of government grants.

Conclusion

The European Peatland Policy Portal introduces a novel, centralised dataset of peatland-affecting policies across administrative levels, departments, geographic areas, and languages, as well as new tools for exploration of that policy dataset through spatial analysis and the power of generative AI. While the portal aims to improve the process of analysing and developing effective, evidence-based peatland policies due to their importance for emissions reduction and climate change mitigation, the tools of spatial analysis of the policy landscape and the RAG interface with the policy collection are transferable to any policy domain. Further, the project methodology including dataset creation can be extended even outside of policy to any domain with scattered, disconnected libraries or repositories needing knowledge integration.

The portal's many tools provide unique analytical capabilities, but their main contribution lies in addressing specific policy issues that scholars and policymakers must address. The link between important policy issues, the related portal technologies, and the policy-relevant results they aim to facilitate is summarised in Table 9.

Policy fragmentation across governance levels is sought to be addressed in the structured policy dataset and its corresponding catalogue page, which consolidates peatland-affecting policies from international, European, national, and subnational levels into a unified, searchable dataset. The portal is intended to lessen the need for dispersed national repositories and ad hoc searches by enabling policymakers to quickly discover all pertinent regulatory instruments impacting peatlands within a specific jurisdiction.

Lack of spatial awareness in policy analysis is sought to be addressed in the Peatland Policy Map, which models policy applicability geographically with administrative extents and peatland locations. The representation aims to allow policymakers to identify regulatory overlaps, gaps, and inconsistencies affecting specific peatland sites, particularly in cross-border or multi-level governance contexts.

Limited transparency and accessibility of complex policy texts is sought to be addressed by the multilingual RAG question-answering tool, enabling users to query the full policy dataset using natural language AI while maintaining explicit citations to authoritative sources. This is intended to support rapid evidence synthesis, improve interpretability, and reduce barriers to policy engagement for both non-legal experts and anyone exploring policies in languages they do not understand.

Looking Forward

The long-term vision of the portal is continual development of new and existing tools as well as updates to the policy and project collections through the portal's suggestion forms. Future projects will include work on the portal as part of the project tasks, allocating time to conduct this development and approval of suggestions. With these added resources, the portal will also conduct evaluations of the efficacy of its tools for policy analysis and development, a task which was outside the scope of the current project.

Future development of the portal tools centres around further AI integration, using it to detect new peatland-affecting policies at national and subnational scales, as well as to update existing policies when they are amended, renewed, or expired. Additionally, the portal aims to incorporate AI into the extraction of relevant policy information like key instruments and metrics in proposed actions, and policy connections such as inheritance between national and subnational implementations of international guidance. These innovations again will be developed in a project focused on peatland policy, but will aim to be transferable to other land use policy domains challenged by fragmented source repositories, multilingual source texts, and departmental siloing.

Additionally, future work includes the development of AI tools for policy analysis tasks more specific than identifying simple policy design elements. Given the need to address climate problems with a robustly diverse array of solutions across departments, the evaluation of policy alignment and divergence is a particularly important step in creating efficient, effective, climate-friendly policies.[41,42] With further development of the RAG tool and knowledge extraction models, AI can help policymakers identify the points of convergence and divergence in policies without taking weeks to read through thousands of pages of text.

We expect this portal will be a source for reference and analysis to improve the creation of evidence-based peatland policies and projects, ultimately leading to healthier environments and emissions reductions throughout Europe and the rest of the world.

A. Supplementary Data

The code of the portal itself and its various tools is open-source and available on GitHub [10] under an MIT license Further documentation of external contributions to the Policy Collection by both partners and stakeholders has also been archived on Zenodo[26]: DOI: 10.5281/zenodo.21376860 (Waskow and Merten, 2026). Partner and Stakeholder Contributions to the PPP's Peatland Policy Collection.

This dataset consists of 3 XLSX files containing policies suggested for the European Peatland Policy Portal's Policy Collection by partners and stakeholders. These files contain A) information on 5 stakeholder workshops held on the portal and its policies (Milestone for project Deliverable E2: 5 national workshops to develop country peatland policy briefs in DE, PL, BE, NL, and IE finished); B) information from a survey sent around to Multi Peat partners concerning the portal and its policy collection (Milestone for project Deliverable E2: Review of Multi-Peat Policy Toolkit involving all partners completed); C) a sheet of policies provided by the German partner NABU.

The data is available under the terms of the Creative Commons Attribution 4.0 International license.

Ethics and Consent

Ethical approval and consent were not required.

Data Availability

The policy data compiled and referenced in this study is available via the European Peatland Policy Portal's Policy Collection [11]. This portal is a dynamic resource; its policy dataset is updated as new policies are published and identified.

For transparency and reproducibility, a static version of the dataset has been archived on Zenodo[25]: DOI: 10.5281/zenodo.21382534 (Ó Brolcháin et al., 2026). The European Peatland Policy Dataset (Version july 2026).

This dataset contains the JSON and XLSX export of the policy collection at the time of analysis and is available under the terms of the Creative Commons Attribution 4.0 International license.

The code of the comparison between Overton and the Peatland Policy Portal is also open-source and available on GitHub [12] under an MIT license.

Further evaluation of the RAG tool has also been archived on Zenodo[35]: DOI: 10.5281/zenodo.21430710 (Khan and Waskow, 2026). Quantitative Evaluations of the PPP RAG Tool (Version v3).

This dataset consists of an XLSX file that contains the RAGAS evaluation of the RAG tool, including the evaluation dataset, the distribution summary of question types, the results summary, and metadata about the model and evaluation. It also includes another XLSX file that includes the comparison of the RAG tool's responses to those of a non-RAG LLM.

The data is available under the terms of the Creative Commons Attribution 4.0 International license.

Acknowledgements

We would like to thank Marcel Šimon, Parisa Zangeneh, Muhammad Imran, Cáit Gleeson, Piaras Ó Giobúin, Samuel Kessler, and the dedicated interns who helped build our portal: Raquel Sabater Canovas, Andrea O'Connor, Lily Bineau, Erica Zanovello, Joseph Casey, Raimey O Boyle, Meret Stursberg, Mella O'Brolchain, Aoife Monaghan, Darren Casserly, Oktavia Dewi Rizka Alam, Shanmathi Magesh, Mary Hachey, David Cunningham, and Sajin Stanislaus.

[10] https://github.com/mawaskow/multipeattools
[11] https://peatlandpolicyportal.eu/policy
[12] https://github.com/marcelsimon39/PPP_v_Overton_Data_Analysis

We additionally would like to thank Jennifer Merten of Naturschutzbund Deutschland e.V., Maria Nolan of Liverpool John Moores, Marelle van der Snoek of Hogeschool Van Hall Larenstein, and Freya Connelly of Abhainn Beo for their expertise.

Flaticon credits for Figure 1: Global, Europe, Belgium, Germany, Ireland, Research, Map, Puzzle, Checklist icons by Freepik; Region icon by Awicon; Retrieve icon by Pop Vectors; Hand shake icon by Aranagraphics.

References

1. Tanneberger F, Tegetmeyer C, Busse S, *et al.*: **The Peatland Map of Europe.** *Mire Peat.* 2017; **19**: 22.
Publisher Full Text

2. Humpenöder F, Karstens K, Lotze-Campen H, *et al.*: **Peatland protection and restoration are key for climate change mitigation.** *Environ. Res. Lett.* 2020; **15**(10): 104093.
Publisher Full Text

3. Urák I, Hartel T, Gallé R, *et al.*: **Worldwide peatland degradations and the related carbon dioxide emissions: the importance of policy regulations.** *Environ. Sci. Pol.* 2017; **69**: 57–64.
Publisher Full Text | Reference Source

4. Chen C, Loft L, Matzdorf B: **Lost in action: Climate friendly use of European peatlands needs coherence and incentive-based policies.** *Environ. Sci. Pol.* 2023; **145**: 104–115.
Publisher Full Text | Reference Source

5. Minasny B, Adetsu DV, Aitkenhead M, *et al.*: **Mapping and monitoring peatland conditions from global to field scale.** *Biogeochemistry.* 2024; **167**(4): 383–425.
Publisher Full Text

6. Nordbeck R, Hogl K: **National peatland strategies in Europe: current status, key themes, and challenges.** *Reg Environ Chang.* 2024; **24**(1): 5.
Publisher Full Text

7. Bonn A, Allott T, Evans M, , *et al.*, editors. *Peatland Restoration and Ecosystem Services: Science, Policy and Practice.* Cambridge: Cambridge University Press; 2016. (Ecological Reviews).
Publisher Full Text

8. Flood K, Wilson D, Renou-Wilson F: **Evidence Synthesis and Knowledge Integration for Sustainable Peatland Management.** *Land.* 2025; **14**(7).
Publisher Full Text | Reference Source

9. Head BW: **Reconsidering evidence-based policy: Key issues and challenges.** *Polic. Soc.* 2010; **29**(2): 77–94.
Publisher Full Text

10. Couwenberg J, Thiele A, Tanneberger F, *et al.*: **Assessing greenhouse gas emissions from peatlands using vegetation as a proxy.** *Hydrobiologia.* 2011; **674**(1): 67–89.
Publisher Full Text

11. Cobb CD: **Geospatial Analysis: A New Window Into Educational Equity, Access, and Opportunity.** *Rev. Res. Educ.* 2020; **44**(1): 97–129.
Publisher Full Text

12. Brody SD, Zahran S, Grover H, *et al.*: **A spatial analysis of local climate change policy in the United States: Risk, stress, and opportunity.** *Landsc. Urban Plan.* 2008; **87**(1): 33–41.
Publisher Full Text | Reference Source

13. Berger T: **Agent-based spatial models applied to agriculture: a simulation tool for technology diffusion, resource use changes and policy analysis.** *Agric. Econ.* 2005; **25**(2–3): 245–260.
Publisher Full Text

14. Macharis C, Pekin E: **Assessing policy measures for the stimulation of intermodal transport: a GIS-based policy analysis.** *J. Transp. Geogr.* 2009; **17**(6): 500–508.
Publisher Full Text | Reference Source

15. Scott G, Rajabifard A: **Sustainable development and geospatial information: a strategic framework for integrating a global policy agenda into national geospatial capabilities.** *Geo-spat. Inf. Sci.* 2017; **20**(2): 59–76.
Publisher Full Text

16. Fezzi C, Bateman IJ: **Structural agricultural land use modeling for spatial agroenvironmental policy analysis.** *Am. J. Agric. Econ.* 2011; **93**(4): 1168–1188.
Publisher Full Text

17. Alsamhi SH, Rowan NJ, Hawbani A, *et al.*: **AI-Empowered Peatland Management for Climate-Resilient and Sustainable Future: Challenges and Opportunities**. Proceedings of the *2024 4th International Conference on Emerging Smart Technologies and Applications (eSmarTA).* IEEE; 2024; 1–8.
Publisher Full Text

18. Firebanks-Quevedo D, Planas J, Buckingham K, *et al.*: **Using machine learning to identify incentives in forestry policy: Towards a new paradigm in policy analysis.** *Forest Policy Econ.* 2022; **134**: 102624.
Publisher Full Text

19. Waskow MA, McCrae JP: **Enhancing Policy Analysis with NLP: A Reproducible Approach to Incentive Classification**. *Proceedings of the 21st Conference on Natural Language Processing (KONVENS 2025): Workshops.* 2025; 74–85.
Reference Source

20. Choi SJ: **Mapping and Comparing Climate Equity Policy Practices Using RAG LLM-Based Semantic Analysis and Recommendation Systems.** *arXiv.* 2026.
Reference Source

21. Nagarajan G, Kumar O, Santhiappan S: **Policy Bot-Reliable Question Answering over Policy Documents.** *arXiv.* 2025.
Reference Source

22. Sun J, Luo Z: **For PKG: A Framework for Constructing Forestry Policy Knowledge Graph and Application Analysis**. Proceedings of the *2025 International Joint Conference on Neural Networks (IJCNN).* IEEE; 2025; 1–8.
Publisher Full Text

23. Belete E: **AI-Powered Climate Policy Analysis Using Natural Language Processing and Large Language Models.** *Zenodo.* 2025.
Publisher Full Text

24. Cantens T: **How will the state think with Chat GPT? The challenges of generative artificial intelligence for public administrations.** *AI Soc.* 2025; **40**(1): 133–144.
Publisher Full Text

25. Brolcháin Ó, Khan MY, Wang Y, *et al.*: **The European Peatland Policy Dataset (Version july 2026) [Dataset].** *Zenodo.* 2026.
Publisher Full Text

26. Waskow MA, Merten J: **Partner and Stakeholder Contributions to the PPP's Peatland Policy Collection [Dataset].** *Zenodo.* 2026.
Publisher Full Text

27. Lourenco M, Fitchett JM, Woodborne S: **Peat definitions: A critical review.** *Progr Phys Geogr Earth Env.* 2023; **47**(4): 506–520.
Publisher Full Text

28. Tegetmeyer C, Kaiser M, Barthelmes A: **The European Wetland Map ('EWM') [Dataset].** *Zenodo.* 2025.
Publisher Full Text

29. Gilet L, Morley TR, Flynn R, *et al.*: **An adaptive mapping framework for the management of peat soils: A new Irish peat soils map.** *Geoderma.* 2024; **447**: 116933.
Publisher Full Text

30. Khan MY, O'Brolchain N, Waskow MA, *et al.*: **Toward Trustworthy Policy Discovery: A Multilingual Hybrid RAG Approach Applied to EU and UK Peatland Restoration**. *Proceedings of the 33rd Irish Conference on Artificial Intelligence and Cognitive Science.* Cham: Springer; 2025.

31. Rawte V, Chakraborty S, Pathak A, *et al.*: **The Troubling Emergence of Hallucination in Large Language Models–An Extensive Definition, Quantification, and Prescriptive Remediations**. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing.* Association for Computational Linguistics. 2023; 2541–2573.
Publisher Full Text

32. Lewis P, Perez E, Piktus A, *et al.*: **Retrieval-augmented generation for knowledge-intensive NLP tasks**. *Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS'20).* Curran Associates Inc; 2020.
Publisher Full Text

33. Song K, Tan X, Qin T, *et al.*: **MPNet: Masked and permuted pre-training for language understanding.** *Adv Neural Inf Process Syst.* 2020; **33**: 16857–16867.
Publisher Full Text

34. Es S, James J, Espinosa-Anke L, *et al.*: **RAGAS: Automated evaluation of retrieval augmented generation**. *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations.* 2024; 150–158.
Publisher Full Text

35. Khan MY, Waskow MA: **Quantitative Evaluations of the PPP RAG Tool (Version v3) [Dataset].** *Zenodo.* 2026.
Publisher Full Text

36. Zhou Y, Liu Y, Li X, *et al.*: **Trustworthiness in retrieval-augmented generation systems: A survey.** *ArXiv.* 2024.
Publisher Full Text

37. Lewis P, Perez E, Piktus A, *et al.*: **Retrieval-augmented generation for knowledge-intensive NLP tasks.** *Adv Neural Inf Process Syst.* 2020; **33**: 9459–9474.
Publisher Full Text

38. Gao Y, Xiong Y, Gao X, *et al.*: **Retrieval-augmented generation for large language models: A survey.** *ArXiv.* 2023.
Publisher Full Text

39. van Belle J, Elferink E: **Site Emissions Tool User Manual: A user-friendly tool to estimate GHG-emission reductions from peat rewetting.** *Interreg North-West Europe Carbon Connects.* **v1**. 2nd; 2020.
Reference Source

40. Ó Brolcháin N, Field C, Tzoulas K, *et al.*: **Validating the bespoke model for the social and economic impact of Carbon farming and developing a theoretical model for the impact of Paludiculture.** *Interreg North-West Europe Care-Peat.* 2023.
Reference Source

41. Greer A, Grant W: **Divergence and continuity after Brexit in agriculture.** *J. Eur. Publ. Policy.* 2023; **30**(11): 2326–2348.
Publisher Full Text

42. Armstrong KA: **Regulatory alignment and divergence after Brexit.** *J. Eur. Publ. Policy.* 2018; **25**(8): 1099–1117.
Publisher Full Text

Open Peer Review

Current Peer Review Status: ✔ ?

Version 2

Reviewer Report 24 August 2026

https://doi.org/10.21956/openreseurope.26597.r78168

© **2026 Das S.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

✔ **Subhankar Das**
Duy Tan University, Da Nang, Vietnam

I am happy with the revisions. I want the authors to add more on implementation, which is optional.

Competing Interests: No competing interests were disclosed.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Version 1

Reviewer Report 20 May 2026

https://doi.org/10.21956/openreseurope.25440.r73732

© **2026 Rabbi M.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

? **Mohammad Fazle Rabbi**
University of Debrecen, Debrecen, Hungary

The manuscript introduces the European Peatland Policy Portal (PPP), a digital infrastructure designed to consolidate scattered and multilingual regulatory frameworks governing peatland restoration across Europe. Developed under the auspices of the EU LIFE Multi Peat project, the toolkit integrates a curated database of 421 policy documents with machine learning and geospatial functionalities. The infrastructure incorporates three primary components: a filterable

tabular policy catalogue, an interactive geospatial map that overlays administrative policy extents onto peat soil distributions, and a retrieval-augmented generation (RAG) question-answering tool designed to navigate cross-lingual regulatory barriers. Additionally, the portal incorporates site-level predictive instruments to evaluate greenhouse gas (GHG) dynamics and economic viability under varying paludiculture scenarios. The stated objective is to mitigate horizontal and vertical fragmentation in environmental governance, thereby facilitating evidence-based land use policymaking in alignment with the EU Nature Restoration Law.

Section-by-Section Critique

Fundamental Significance and Novelty

The manuscript addresses a documented administrative barrier in environmental governance, specifically the fragmentation of cross-jurisdictional environmental directives. Compiling a multilingual repository containing over 400 documents across 27 nations represents a pragmatic resource contribution for public administration practitioners. However, the conceptual novelty of the study remains poorly articulated within the broader context of information science and digital governance. The translation of established natural language processing architectures, such as retrieval-augmented generation, into a specific environmental sub-domain constitutes an incremental application rather than a methodological paradigm shift. The theoretical positioning of how spatial policy mapping alters governance networks or decision-making efficiency relies on speculative assertions rather than validated administrative frameworks.
Methodological Rigor and Reproducibility
The data collection framework exhibits significant methodological vulnerabilities that undermine the reproducibility of the dataset. The execution of Phase I and Phase II acquisitions relied heavily on unstructured Google searches and unstandardized national workshops. The deployment of generic keyword queries (e.g., "peatland", "biodiversity") lacks systematic operational boundaries. The reliance on Google Translate for the initial filtration of non-English legislative texts introduces systematic error, as legal terminology requires precise taxonomic cross-mapping rather than automated algorithmic translation. Furthermore, the inclusion criterion that permits documents lacking explicit nomenclature if they "potentially have an impact on peatlands" introduces substantial investigator bias. The absence of inter-rater reliability metrics or formal algorithmic protocols for document validation prevents independent verification of the corpus.

Data Integrity and Statistical Accuracy

The empirical presentation exhibits internal contradictions and structural omissions. The narrative text states that the dataset encompasses 421 policy documents across 27 countries. However, the preceding text detailing the conclusion of Phase II specifies a total of 405 policies. The manuscript fails to provide a quantitative breakdown or chronological reconciliation of the 16 additional documents ostensibly captured during Phase III. Figure 3 displays quantitative outputs regarding policy counts by country and language percentage, yet these data are presented as a descriptive static visualization without an underlying statistical cross-tabulation matrix. No statistical validation or error-bound modelling is provided for the integration of the European Wetland Map (EWM) polygons, despite the acknowledged structural heterogeneity and variable definitions across national boundary datasets.

Interpretive Logic and Discussion

The discussion section shifts toward descriptive system documentation rather than analytical critique. The evaluation of system efficacy relies on a highly restricted qualitative sample consisting of two illustrative queries. The dismissal of a failed retrieval event (Question 1) as merely "off-topic" demonstrates a lack of critical technical evaluation. Instead of quantifying system accuracy, precision, or recall through established computational metrics (e.g., ROUGE, BLEU, or human expert panel Likert scoring), the authors describe an ad-hoc modification termed a "novel fact injection method". This mechanism is introduced without technical architecture specifications or formal comparative validation against the baseline RAG pipeline. The comparative synthesis against existing platforms (e.g., EUR-Lex, Carbon Gap) remains superficial, focusing on database scope rather than presenting a rigorous benchmark of retrieval performance or user task efficiency.

Categorization of Critique

Major Flaws
• Major Flaw 1: Absence of Quantitative System Evaluation The evaluation of the retrieval-augmented generation toolkit is structurally deficient, relying on two non-reproducible example queries rather than a validated empirical test suite. To establish suitability for publication, the system must be rigorously benchmarked using quantified validation metrics. The investigators must compute precision, recall, and context-relevance scores across a standardized evaluation dataset ($N \geq 100$ queries) evaluated by independent legal and environmental experts.

• Major Flaw 2: Non-Systematic and Non-Reproducible Data Acquisition Methodology The data collection methodology introduces unquantified selection bias through its reliance on opportunistic web search queries, unvetted national workshops, and automated machine translation tools. The lack of a formal, algorithmic search protocol or formalized legal taxonomy for multi-jurisdictional policy identification compromises the integrity of the database. A reproducible verification protocol must be instituted to standardize document inclusion across all 27 targeted jurisdictions.

• Major Flaw 3: Structural Numerical Discrepancy in Dataset Reporting A contradiction exists within the reporting of the core dataset dimensions. The manuscript states that Phase II concluded with 405 policy documents. The subsequent section asserts a total repository size of 421 documents. The text lacks a systematic statistical accounting of the 16 additional policies, rendering the data description internally inconsistent.

Minor Flaws
• Minor Flaw 1: Unsubstantiated Inclusions via Subjective Criteria The guideline allowing the incorporation of policy documents that omit core terms but possess a speculative "potential impact" introduces subjective investigator bias. This boundary condition must be constrained by explicit thematic or legal proxy indicators.

• Minor Flaw 2: Improper Mathematical Notation Placement Statistical parameters and quantitative dimensions within the descriptive narrative are not formatted consistently in standard mathematical text. The specification of the land area percentage (5.85%) and the associated absolute spatial area (241,213 square km) must be rendered uniformly in mathematical notation.

• Minor Flaw 3: Omission of Underlying Data Tables for Visualizations Figure 3 and Figure 4 present data regarding country counts, linguistic distributions, and governance levels, but the manuscript lacks corresponding tabular matrices detailing the exact frequencies and percentages. These data must be provided as supplementary tables to ensure data transparency.

Final Recommendation
Major Revision

The manuscript presents a functional digital infrastructure that addresses a clear practical challenge in European environmental policy coordination. However, the submission currently resembles a technical system documentation report rather than a peer-reviewed research article. The scientific merit of the paper is limited by the non-systematic nature of the data collection protocol and a total absence of quantitative validation for the machine learning retrieval pipeline. To achieve publication standard, the authors must execute a rigorous, quantitative evaluation of the RAG tool performance using established computational and human-expert metrics. Furthermore, the methodological description must be revised to clarify inclusion parameters and resolve numerical discrepancies across the text sections. The core architectural framework holds merit, but empirical rigor must be established to substantiate the utility of the portal.

Is the work original in terms of material and argument?
Partly

Does it sufficiently engage with relevant methodologies and secondary literature on the topic?
Partly

Is the work clearly and cogently presented?
Partly

Is the argument persuasive and supported by evidence?
No

If any, are all the source data and materials underlying the results available?
Partly

Does the research article contribute to the cultural, historical, social understanding of the field?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: machine learning, sustainability, environmental sustainability, renewable energy

I confirm that I have read this submission and believe that I have an appropriate level of

expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 16 Jun 2026

Margaret Waskow

Hello, this is just to say that I have contacted the relevant authors regarding necessary changes, experiments, and further documentation. We will submit a revised version of the article once those changes and additional information are returned. Thank you for your review.

Competing Interests: No competing interests were disclosed.

Author Response 23 Jul 2026

Margaret Waskow

Thank you again for your review. In our second draft, we have attempted to address the following points. As the authors do not yet have the final typeset version, references to the update manuscript are provided by section and paragraph if necessary, or reproduced in the response if the addition was short enough.

Reviewer #2, Comment #1: Absence of Quantitative System Evaluation.
Author Response: Thank you for this valuable recommendation. We agree that a rigorous evaluation was required.
Author Action: We updated the manuscript by: A) performing a RAGAS evaluation of the RAG tool on 58 question and source-document pairs covering multiple languages and governance levels; B) adding a baseline comparison with a non-RAG LLM. The updates in the revised version of the manuscript can be found in the "Policy Exploration Tools > RAG Question-Answering" section's new subsections "Quantitative Evaluation: RAGAS" and "Quantitative Evaluation: Baseline Comparison with non-RAG LLM."

Author Note: The metrics reported in the RAGAS evaluation are context precision (overall: 0.895), answer relevancy (overall: 0.638), and citation precision (overall: 1.000)-- classical IR metrics (recall, precision) are not applicable at scale here as per-query relevance labels do not exist for this corpus; context precision serves as the accepted proxy. Reviewer 1 requested at least 30 examples as opposed to Reviwer 2's 100, so the authors compromised for a total of 58 queries across various topics. Additionally, a formal hallucination audit by independent raters has not yet been conducted due to a lack of resources, however the RAGAS framework targets hallucinations in its analysis and is the standard for RAG system evaluation.

Reviewer #2, Comment #2: Non-Systematic and Non-Reproducible Data Acquisition Methodology.
Author Response: We appreciate this observation of the original methodology.
Author Action: We updated the manuscript to identify and acknowledge the presence of selection bias and investigator bias in the first two phases of policy acquisition, then

articulated the new criteria and protocols as of Phase III to help address the bias. The first exploration of the bias can be found in the first paragraph of the "The European Peatland Policy Dataset > Dataset Curation" section where the following text was added: "As guided by the Multi Peat project deliverables, the peatland policy collection's construction included not only web and database searches but also collaborations with project partners and stakeholder workshops. While the inclusion of policies suggested by experts introduces selection bias into the policy collection process, translation challenges and the decentralised nature of many national and sub-national policy repositories left gaps in the search method's results which experts of the country's policy were able to fill. Thus, as it stands, the policy collection is not presented as a necessarily comprehensive or exhaustive list of the peatland-affecting policies of Europe, but rather as an ongoing attempt to consolidate these policies based on documented criteria and expert input." The second exploration of the bias can be found in the third paragraph of the "The European Peatland Policy Dataset > Dataset Curation > Phase II" section where the following text was added: "For policies recommended by experts, the source of the recommendation was taken as evidence of the policy's relevance, so some policies were included which did not contain the three core keywords; this exceptionalism introduced investigator bias into the dataset which then has been addressed in Phase III." The description of how these issues have been addressed as of the revised Phase III can be found in the "The European Peatland Policy Dataset > Dataset Curation > Phase III" section where the following text was added: "At this point, researchers are focusing more on regional and local policies, requiring the identification and utilisation of siloed local authority and departmental policy repositories. Additionally, the analysts returned to policies from earlier phases to confirm their inclusion or validate their exclusion; after reviewing the expert-recommended policies which were relevant to peatlands but lacked the three core keywords, researchers identified four more keywords indicating a policy's non-explicit relevance to peatlands. This expanded the inclusion criteria as of Phase III, requiring policies to either contain the linguistic equivalents of explicit peatland relevance (i.e. "Peat", "Peatland", "Wetland"), or the linguistic equivalents of non-explicit peatland relevance (i.e. "Organic soil", "Carbon soil", "Water table", "Turf"). This criteria update allowed the inclusion of some formerly rejected policies found via search, and enabled the researchers to address some of the investigator bias introduced by previously holding expert-suggested policies to a separate standard. Policies collected in this phase additionally come from workshops held with experts in different European countries." Additionally, we updated the manuscript by adding further information about the dataset curation, including: A) a new version of the policy dataset which contains the date of first addition to our internal policy list and most recent publication of the data entry onto the portal, in order to facilitate better understanding of what policies were collected when/which phase; B) regarding the expert consultation, a supplementary dataset documenting policies gathered through partners and stakeholders; C) updated inclusion/exclusion criteria as of Phase III in Table 2; D) a new table (Table 3) showing 4 worked examples of borderline inclusion and exclusion decisions. Updates A and B can be accessed via Zenodo, with links found in the "Data Availability" and "Supplementary Data" sections, and discussion of the datasets found in the "The European Peatland Policy Dataset > Dataset Curation" section (paragraph 1). Update C can be found in Table 2 in the "The European Peatland Policy Dataset > Dataset Curation > Inclusion Criteria" section. Update D can be found in Table 3 in the "The European Peatland Policy Dataset > Dataset Curation > Edge Cases" section.

Reviewer #2, Comment #3: Structural Numerical Discrepancy in Dataset Reporting.
Author Response: Thank you for identifying this ambiguity.
Author Action: We have updated the underlying dataset of policies on Zenodo to include dates of the policy's first discovery by the researchers (creation date) and most recent publishing of the data entry to the live site (publication date). The authors believe that this will help to address the ambiguity of total policy count from later additions to the dataset in Phase III.

Reviewer #2, Comment #4: Unsubstantiated Inclusions via Subjective Criteria.
Author Response: Thank you for this suggestion.
Author Action: We updated the manuscript by acknowledging the presence of investigator bias in Phase II's inclusion of policy documents that omit the 3 original inclusion keywords but possess speculative "potential impact." Additionally, we describe Phase III's reworking of the inclusion keywords to add 4 keywords of indirect impact on peatlands, which then were applied retroactively to the policy collection to remove policies which did not contain at least one of the core 7 keywords, as well as to include formerly rejected policies which did include one of the new 4 keywords. The full explanation of changes and their location in the text can be found in the above response to Comment #2.

Reviewer #2, Comment #5: Improper Mathematical Notation Placement.
Author Response: Thank you for this feedback.
Author Action: The authors were uncertain whether they understood this issue correctly, but they will raise it with the ORE typesetter, have moved the values next to each other in the sentence for clearer understanding, and further have replaced the calculated EU values with the reported European region values. If the authors did not address this issue correctly in the second draft, they request that Reviewer 2 clarify the issue so they can correct it. This change can be found in the second sentence of the Introduction.

Reviewer #2, Comment #6: Omission of Underlying Data Tables for Visualizations.
Author Response: Thank you for improving transparency.
Author Action: We updated the manuscript by updating figures to display exact underlying frequencies directly instead of relative percentage composition, eliminating ambiguity. This change can be found in Figures 3 and 4 in the "The European Peatland Policy Dataset > Characteristics" section.

Competing Interests: No competing interests were disclosed.

Reviewer Report 20 May 2026

https://doi.org/10.21956/openreseurope.25440.r73736

© **2026 Das S.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Subhankar Das
Duy Tan University, Da Nang, Vietnam

The following revisions are required before the manuscript can be reconsidered. Items 1 through 4 are non-negotiable for indexing; items 5 through 9 substantially strengthen the contribution.

1. **Provide a real evaluation of the RAG tool.** Construct a held-out evaluation set of at least 30 question and source-document pairs covering multiple languages and governance levels. Report retrieval precision and recall, a hallucination audit by at least two raters, and a baseline comparison against a non-RAG LLM on the same questions.
2. **Add a technical specification of the RAG system.** In the main text or an appendix, document the embedding model, chunking strategy, hybrid retrieval weighting, LLM and generation parameters, multilingual handling, and citation grounding mechanism in sufficient detail to permit replication without reference to Khan et al. (2025).
3. **Reframe the dataset characteristics.** Clearly distinguish between (a) properties of the dataset as constructed and (b) inferences about European peatland governance. The English over-representation and the national-level skew are artefacts of methodology and should be presented as such, not as findings.
4. **Recalibrate efficacy claims throughout.** Audit the abstract, introduction, and conclusion for verbs such as enables, supports, addresses, assists. Either reserve these for capabilities that have been measured or replace with descriptive language (is designed to, is intended to). The current draft frequently writes design intentions as if they were demonstrated outcomes.
5. **Document the dataset curation protocol rigorously.** Resolve the duplicate methane strategy entry in Table 1. Document the expert consultation behind the 19 keywords. Provide at least three worked examples of borderline inclusion and exclusion decisions. Report a recall check using known peatland policies from two or three non-English countries.
6. **Report categorisation reliability.** Either re-code a sample of at least 50 policies with a second annotator and report Cohen's kappa, or be explicit that the category labels in Figure 4 are single-coder and should be read accordingly.
7. **Address spatial false positives in the policy map.** Identify which policies in the dataset have subject-matter sub-scoping (e.g. Natura 2000 only, organic soils only) that the administrative-boundary polygon flattens. Either implement sub-scoping in the metadata or transparently flag the limitation in the manuscript and the interface.
8. **Validate or qualify the peatland polygons.** Report a small ground-truth check on the peatland polygons in two or three countries, or restate the polygons as approximate land-cover indicators rather than authoritative peatland maps.
9. **Bound the transferability claim.** Replace the universal claim with a conditional one: specify the properties of a domain (fragmented sources, multilingual texts, ambiguous categorisation) under which the methodology is expected to transfer.

The European Peatland Policy Portal is a substantive piece of infrastructure for European environmental policy research. The dataset is genuinely useful, the spatial policy retrieval is novel, and the RAG interface is a defensible response to a real multilingual access problem. The manuscript in its current form, however, conflates building a system with demonstrating its value, and defers the technical and evaluative scrutiny that an indexed venue requires.

Is the work original in terms of material and argument?

Partly

Does it sufficiently engage with relevant methodologies and secondary literature on the topic?
No

Is the work clearly and cogently presented?
Partly

Is the argument persuasive and supported by evidence?
Partly

If any, are all the source data and materials underlying the results available?
Partly

Does the research article contribute to the cultural, historical, social understanding of the field?
Yes

Competing Interests: No competing interests were disclosed.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 16 Jun 2026

Margaret Waskow

Hello, this is just to say that I have contacted the relevant authors regarding necessary changes, experiments, and further documentation. We will submit a revised version of the article once those changes and additional information are returned. Thank you for your review.

Competing Interests: No competing interests were disclosed.

Author Response 23 Jul 2026

Margaret Waskow

Thank you again for your review. In our second draft, we have attempted to address the following points. As the authors do not yet have the final typeset version, references to the update manuscript are provided by section and paragraph if necessary or reproduced in the response if the addition was short enough.

Reviewer #1, Comment #1: Provide a real evaluation of the RAG tool.

Author Response: Thank you for this valuable recommendation. We agree that a rigorous evaluation was required.
Author Action: We updated the manuscript by: A) performing a RAGAS evaluation of the RAG tool on 58 question and source-document pairs covering multiple languages and governance levels; B) adding a baseline comparison with a non-RAG LLM. The updates in the revised version of the manuscript can be found in the "Policy Exploration Tools > RAG Question-Answering" section's new subsections "Quantitative Evaluation: RAGAS" and "Quantitative Evaluation: Baseline Comparison with non-RAG LLM."
Author Note: The metrics reported in the RAGAS evaluation are context precision (overall: 0.895), answer relevancy (overall: 0.638), and citation precision (overall: 1.000)-- classical IR metrics (recall, precision) are not applicable at scale here as per-query relevance labels do not exist for this corpus; context precision serves as the accepted proxy. Additionally, a formal hallucination audit by independent raters has not yet been conducted due to a lack of resources; however the RAGAS framework targets hallucinations in its analysis and is the standard for RAG system evaluation.

Reviewer #1, Comment #2: Add a technical specification of the RAG system.

Author Response: Thank you for highlighting the need for reproducibility.
Author Action: We updated the manuscript by adding relevant technical specifications to the RAG tool section, namely: Document ingestion and chunking, Embedding model, Hybrid retrieval, Metadata filtering, Multilingual handling, Generation parameters, and Citation grounding. The updates in the revised version of the manuscript can be found in the "Policy Exploration Tools > RAG Question-Answering" section's new subsection "RAG System Architecture."

Reviewer #1, Comment #3: Reframe the dataset characteristics.
Author Response: Thank you for identifying the difference between dataset properties and governance inferences—this is a very important distinction to make.
Author Action: We updated the manuscript by rewriting the section to describe only the constructed dataset characteristics and removed unsupported governance interpretations as shown in the revised manuscript's "The European Peatland Policy Dataset > Characteristics" section and its subsections "Language Composition", "Country Composition", and "Governance Level and Category Composition. "

Reviewer #1, Comment #4: Recalibrate efficacy claims throughout.
Author Response: We agree that some wording overstated validated performance.
Author Action: We updated the manuscript by replacing outcome-oriented language with design-oriented wording throughout the manuscript as shown in the revised manuscript's Abstract, Introduction, Conclusion, and in the sections describing the tools and their purposes. A small selection of the updates is reproduced below in the format of Section Title: "original text" > "revised text." Abstract: "By enabling systematic policy examination and analysis...the portal assists the coordination and efficiency gaps of European peatland governance." > "**In its drive to support** systematic policy examination and analysis...the portal **aims to bridge** efficiency gaps and **aid the coordination** of European peatland governance." Abstract: "While this portal demonstrates the efficacy of these tools for peatland policy specifically, its framework is transferable..." > "While the portal **proposes**

these tools for peatland policy specifically, **the authors argue that** the framework is transferable..." Introduction > The European Peatland Policy Portal (paragraph 2): "...we highlight how high-resolution, data-driven insights "> "...**the authors argue that** high-resolution, data-driven insights..." The Project Perspective (paragraph 2): "...the portal addresses policy and knowledge gaps..." > "...the portal **aims to address** policy and knowledge gaps..." Policy Exploration Tools > Policy Collection (paragraph 2): "...dataset in this format to streamline the most common policy discovery tasks..." > "...dataset in this format **with the goal of streamlining** the most common policy discovery tasks..." Conclusion (paragraph 2): "...the policy-relevant results they facilitate..." > "...the policy-relevant results they **aim to** facilitate..."

Reviewer #1, Comment #5: Document the dataset curation protocol rigorously.
Author Response: Thank you for this important suggestion.
Author Action: We updated the manuscript by adding additional information about the dataset curation, including: A) a new version of the policy dataset which contains the date of first addition to our internal policy list and most recent publication of the data entry onto the portal, in order to facilitate better understanding of what policies were collected when/which phase; B) regarding the expert consultation, a supplementary dataset documenting policies gathered through partners and stakeholders; C) a new table (Table 3) showing 4 worked examples of borderline inclusion and exclusion decisions; D) the duplicate “methane strategy” entry in Table 1 replaced with the “Greenhouse gas emission/eu ets” entry it was meant to have—thank you again for catching this; E) updated inclusion/exclusion criteria as of Phase III in Table 2. Updates A and B can be accessed via Zenodo, with links found in the "Data Availability" and "Supplementary Data" sections, and discussion of the datasets found in the "The European Peatland Policy Dataset > Dataset Curation" section (paragraph 2). Update C can be found in Table 3 in the "The European Peatland Policy Dataset > Dataset Curation > Edge Cases" section. Update D can be found in Table 1 in the "The European Peatland Policy Dataset > Dataset Curation > Phase II" section. Update E can be found in Table 2 in the "The European Peatland Policy Dataset > Dataset Curation > Inclusion Criteria" section.
Author Note: The authors unfortunately did not understand the last request to report a recall check using known peatland policies from two or three non-English countries, however they believe they addressed the spirit of the request through the new dataset of policies gathered from policy experts in different countries.

Reviewer #1, Comment #6: Report categorisation reliability.
Author Response: We appreciate this recommendation.
Author Action: We updated the manuscript by revising the discussion of Figure 4 to explicitly state that the category labels are single-coder and should be read accordingly. The updates can be found in the "The European Peatland Policy Dataset > Dataset Curation > Categories" section where this sentence was added: "These categories are assigned by single coders for the purpose of providing general context for the policy topic, not for the purpose of analysis." An additional update addressing this can also be found in the "The European Peatland Policy Dataset > Characteristics > Governance Level and Category Composition" section where this sentence was added: "The categories assigned to the policies are single-coder, so they should be taken as a general indication of policy topic instead as a reliable basis for analysis."

Reviewer #1, Comment #7: Address spatial false positives in the policy map.
Author Response: Thank you for identifying this limitation.
Author Action: We updated the manuscript by flagging the spatial false positives introduced by policies with subject-matter sub-scoping—this was another greatly appreciated highlight from Reviewer 1. The update can be found in the "Policy Exploration Tools > Peatland Policy Map > The Policy Polygons" section where the following paragraph was added: "There are some limitations to assigning policies spatial boundaries by administrative extent, as some policies target only specific sites (e.g. Natura 2000 sites) or specific land types (e.g. organic soils) within a region. For example, the Dutch Natura 2000 Management Plan for Witte Veen (Natura 2000 Beheerplan Witte Veen) only affects the specific Witte Veen site, but the policy is assigned the same administrative boundaries as the publisher, the Province of Overijssel; when a user clicks anywhere within the Province of Overijssel, even outside the boundaries of Witte Veen, this policy will be retrieved and listed as affecting that location. The flattening of policy sub-scoping through use of general administrative boundaries to define spatial effect does introduce false positives to the retrieved policy lists, so users are alerted to this limitation in the Policy Map's explanation page [7]."

Reviewer #1, Comment #8: Validate or qualify the peatland polygons.
Author Response: We agree that the limitations should be transparent.
Author Action: We updated the manuscript by acknowledging the limitations of the peatland polygons and their lack of validation, as well as by restating them as approximations for context instead of authoritative maps for analysis. The update can be found in the "Policy Exploration Tools > Peatland Policy Map > The Peatland Polygons" section where the following paragraph was added: "The display of the peatland polygons in their original, full-resolution form caused latency issues in the map due the sheer amount of data being loaded by the server. In order to address this, scale dependencies were introduced so that very small polygons were hidden and the rest of the polygons were simplified when zoomed out on the map. These changes, and the lack of in-house validation of the polygons through ground-truthing, mean that the peatlands displayed on the map should not be taken as authoritative peatland maps, and instead as approximate land-cover indicators for context in understanding the peatland policy landscape."

Reviewer #1, Comment #9: Bound the transferability claim.
Author Response: Thank you for this suggestion.
Author Action: We updated the manuscript by rephrasing the transferability claim to limit it to land use policy domains challenged by fragmented source repositories, multilingual source texts, and departmental siloing. This update can be found in the "Conclusion > Looking Forward" section (paragraph 2): "These innovations again will be developed in a project focused on peatland policy, but will **aim to be transferable to other land use policy domains challenged by fragmented source repositories, multilingual source texts, and departmental siloing.**"

Competing Interests: No competing interests were disclosed.