FROM *POSTFACTUM* TO *ANTEFACTUM* GEOGRAPHY: WHAT ARTIFICIAL INTELLIGENCE WOULD HAVE CHANGED IN SEVEN EARLIER STUDIES

A Retrospective Human–AI Experiment in Technical Geography

Ionel HAIDU[1,2]

A Letter from the Editor

DOI: 10.21163/GT_2026.212.12

ABSTRACT

Discussions of artificial intelligence in Geography frequently take the form of inventories of possible applications. This Letter adopts a different strategy. Instead of asking what AI might theoretically do, it asks what contemporary AI would have changed in seven geographical studies previously authored or co-authored by the editor and completed without access to present-day AI technologies. The studies concern forest-cover change, storm-induced windthrows, floodplain delineation, urban imperviousness and runoff, trend-signal detection, road-accident distribution, and atmospheric pollution during the COVID-19 lockdown. Each paper is examined through the same retrospective protocol: the original geographical question and technical workflow are reconstructed; possible AI entry points are identified; the potential gains in speed, accuracy, uncertainty assessment, knowledge discovery, prediction and decision support are evaluated; and the data and validation requirements of the proposed redesign are discussed. The exercise does not claim that AI would necessarily have produced superior empirical results, because the original datasets were not reprocessed. It demonstrates, however, that AI could have enlarged the range of questions that could be formulated and tested. The seven cases show a common transition from fixed workflows to adaptive learning, from deterministic maps to probability fields, from separate datasets to multimodal integration, and from *postfactum* explanation to *antefactum* prediction, counterfactual analysis and intervention design. AI is therefore proposed not as a substitute for geographical reasoning, but as the learning layer of Technical Geography.

Keywords: *GeoAI; spatial reasoning; hybrid modelling; uncertainty; counterfactual analysis; decision support; human–machine collaboration*

1. TECHNICAL GEOGRAPHY BEFORE AND WITH ARTIFICIAL INTELLIGENCE

In 2016, I attempted to define Technical Geography as an emerging field grounded in the spatial and temporal attributes of geographical processes, particularly autocorrelation and frequency, and supported by spatial techniques, statistical methods and modern technologies. I argued that digital geographical information, remote sensing, GIS, modelling and advanced data processing were not peripheral tools but necessary components of contemporary geographical research. I also anticipated that the technical component would progressively become present in all fields of Geography (Haidu, 2016).

In 2024, I developed this argument further by suggesting that Geography should move from a predominantly *postfactum* science, concerned with phenomena that have already occurred, towards an *antefactum* science capable of anticipation, prediction and proactive intervention. Technical methods, spatial modelling, GIS, remote sensing, LiDAR, UAV observations and statistical analysis were presented as the foundations of this transition (Haidu, 2024).

[1] *LOTERR (Laboratoire d'observations des territoires), Université de Lorraine, F-57000 Metz, France. ionel.haidu@univ-lorraine.fr*

[2] *STAR-UBB (Scientific and Technological Advanced Research Institute), Babeş-Bolyai University, 400084 Cluj-Napoca, Romania.*

Artificial intelligence now requires a further extension of both propositions. AI should not be regarded merely as one additional technology placed alongside GIS, remote sensing or spatial statistics. Its distinctive contribution is the capacity to learn relationships from examples, to update those relationships when new observations become available, to integrate heterogeneous sources of information, to explore large numbers of alternative scenarios and to quantify the uncertainty attached to a prediction.

Classical technical modelling normally begins with a structure specified by the researcher: variables are selected, equations are defined, thresholds are established and the model is then applied to observations. In AI-assisted research, part of this structure can be learned from data. The researcher continues to define the scientific problem, the spatial and temporal support, the relevant processes, the admissible variables and the conditions of validation, but the model may discover nonlinear relationships or interactions that were not explicitly introduced beforehand.

GeoAI has been described as a spatially explicit form of artificial intelligence directed towards geographical knowledge discovery (Janowicz et al., 2020). This definition is important. A generic AI model may identify statistical regularities, whereas a geographical AI model must determine whether those regularities remain meaningful when location, scale, neighbourhood, spatial dependence, temporal evolution and geographical process are considered. Recent Earth-system research has similarly argued for hybrid approaches in which data-driven learning and process understanding reinforce one another rather than compete (Reichstein et al., 2019). Theory-guided data science makes the same point in more general terms: scientific consistency and domain knowledge should enter the learning process, particularly where observations are limited and physical processes are complex (Karpatne et al., 2017).

Consequently, the development of AI does not reduce the need for geographical reasoning. It increases it. The larger the space of patterns and models that a machine can generate, the more necessary it becomes for the geographer to decide which relations are spatially meaningful, physically plausible, causally defensible and socially relevant.

2. THE INVERTED RETROSPECTIVE PARADIGM

A conventional Letter on AI in Geography could enumerate a large number of possibilities: automatic image classification, object detection, prediction of environmental hazards, processing of geographical texts, optimisation of transport networks, climate modelling or decision support. Such a catalogue would be informative, but it would remain largely prospective and theoretical.

The present Letter adopts an inverted retrospective paradigm. It begins not with the capabilities of AI, but with completed geographical studies. It asks a counterfactual question: How might the research design, processing workflow, results and conclusions of an earlier geographical study have changed if contemporary AI methods had been available when the study was conducted?

Seven studies in which I was an author or co-author were deliberately selected. Restricting the corpus to my own work has two purposes. First, it avoids retrospectively criticising other researchers for not having used methods that were unavailable at the time of their research. Second, it transforms the exercise into methodological self-examination. The earlier papers are not presented as obsolete or inadequate. They are treated as authentic records of the technical possibilities, data constraints and scientific reasoning available at particular moments.

Each original study therefore functions as a methodological baseline. Its aims, datasets, models, results and acknowledged limitations are retained. The AI-assisted redesign does not replace this baseline; it asks what additional pathways might have become possible. In some cases, AI would probably have accelerated an operation already performed. In others, it might have allowed uncertainty to be quantified, different data streams to be fused, a previously inaccessible relation to be tested, or a retrospective analysis to become predictive.

This distinction also imposes an important epistemological limit. Because the original datasets have not been reprocessed with the proposed AI models, the present Letter cannot claim that the resulting maps, coefficients or forecasts would certainly have been more accurate. Expressions such as AI could have enabled, might have revealed, or would probably have reduced are scientifically more appropriate than unconditional assertions of superiority. The experiment is not an empirical replication of the seven studies. It is an experiment in research design: a systematic comparison between the workflow that was historically implemented and the expanded workflow that can now be conceived.

3. METHOD OF THE HUMAN–AI RETROSPECTIVE EXPERIMENT

The seven studies cover different components of geographical research: forest remote sensing, severe-weather analysis, hydraulic modelling, urban hydrology, time-series analysis, transport geography and atmospheric pollution. This disciplinary diversity makes it possible to evaluate whether AI contributes only to a narrow class of computational problems or whether it modifies geographical reasoning more broadly (**Table 1**).

Table 1.

Common protocol applied to the seven retrospective cases.

Retrospective question	Analytical purpose
What geographical question was originally addressed, and through what technical workflow?	Reconstruct the historical methodological baseline.
At which stages could contemporary AI enter the workflow?	Identify concrete rather than generic AI applications.
What type of benefit might result?	Distinguish speed, accuracy, integration, uncertainty, discovery, prediction and decision support.
What data and validation would be required?	Avoid presenting a plausible AI proposal as an already demonstrated result.

The potential contribution of AI is classified into six categories: workflow acceleration, improved extraction or classification, integration of heterogeneous data, uncertainty quantification, discovery of additional relationships, and transformation from retrospective analysis into prediction or decision support.

An AI proposal is retained only when it is compatible with the geographical process and the data structure of the original study. A technically possible model is not necessarily geographically justified. For example, a high predictive score is insufficient if the model ignores spatial dependence, uses information from the future during training, confuses exposure with risk, or violates a physical conservation law.

The retrospective assessment reported below was developed through a documented, iterative dialogue between the author and an artificial-intelligence system. The author defined the conceptual framework, selected the case studies and evaluated the scientific relevance of the proposed methodological alternatives. The complete disclosure of the respective human and AI contributions is provided at the end of the Letter.

All seven papers were already published; no confidential third-party manuscript was used. The exercise also separates AI-assisted reasoning from scientific evidence: the proposed methods are supported by published methodological literature, while the AI output itself is not treated as a bibliographical source.

4. FROM THREE FOREST-COVER MAPS TO LEARNING FOREST TRAJECTORIES

Original research design. The first study combined forestry maps and inventory tables with Landsat images acquired in 1988, 2000 and 2010. A common spectral library was constructed, and the Spectral Angle Mapper (SAM) classifier was applied to identify forest and non-forest land-cover classes. The resulting maps were processed in GIS to assess deforestation in three Apuseni Mountain catchments, with forest-cover changes approaching 22% of one catchment (Costea et al., 2012). The method had an important strength: it preserved spectral interpretability and applied a common classification logic to the three images. Its validation produced high overall accuracy and Kappa values of 0.98, 0.94 and 0.91. Nevertheless, the analysis was based on three temporal snapshots and identified difficulties associated with mixed forests, the water-stress conditions visible in the 2000 image, vegetation water content, illumination and the precision of forestry inventories.

AI-assisted redesign. I would retain SAM as a physically interpretable benchmark. AI should not remove a transparent method merely because a more complex method is available. I would add comparative machine-learning models, particularly Random Forests (Breiman, 2001), using not only spectral bands but also vegetation indices, altitude, slope, aspect and spatial texture. The more substantial change would be to formulate the problem directly as multitemporal change detection. Fully convolutional Siamese architectures can analyse two co-registered images simultaneously and learn the spatial and spectral characteristics of change (Daudt et al., 2018), while a U-Net-type encoder-decoder can preserve both contextual information and local boundaries (Ronneberger et al., 2015). Instead of using only three dates, the complete available Landsat archive could be introduced into a temporal model. It could learn different trajectories: stable forest, gradual degradation, abrupt removal, temporary drought stress, regeneration and afforestation. Training could be organised through active learning, with the model identifying the spatial units for which expert labelling would reduce uncertainty most efficiently.

Potential added value and validation requirements. The output would no longer consist only of land-cover maps and differences between three dates. It could include the probable date of change, the probability that a change represents deforestation rather than temporary stress, the rate of vegetation recovery and the uncertainty associated with each mapped object. The spatial pattern of forest loss could then be connected to a hydrological model or to an AI emulator of that model. The study could ask not only how much forest disappeared, but which spatial configurations of forest removal most strongly affect runoff and flash-flood generation. A rigorous evaluation would require temporally independent reference samples, harmonisation among Landsat sensors and forestry information capable of distinguishing harvesting, degradation, drought damage and subsequent recovery. Remote sensing and AI should not be expected to infer the legal or socioeconomic cause of every change without additional evidence. The principal contribution would be a change from comparing maps to learning geographical trajectories.

5. FROM RECONSTRUCTING A WINDTHROW EVENT TO PREDICTING ITS SPATIAL IMPACT

Original research design. The windthrow study examined the severe event of 20 July 2011 in the Apuseni Mountains. Areas affected by windthrows were detected through Landsat imagery and differences in NDVI, while synoptic maps and WSR-98D Doppler-radar products were used to reconstruct the regional and mesoscale atmospheric evolution. The radar analysis included reflectivity, storm movement, Vertically Integrated Liquid water and echo-top height (Furtuna et al., 2018). The study identified convective descending movements, three storm cores and short-duration wind gusts exceeding 23 m s^{-1}. It concluded by asking whether recurrent storm trajectories and preferential travel corridors could be identified in relation to relief and baric configuration.

AI-assisted redesign. The first AI component would be a semantic segmentation model trained on pre- and post-event satellite images to delineate windthrow patches from spectral, textural and morphological evidence rather than from a single NDVI-difference criterion. The second component would operate on the complete sequence of radar images. ConvLSTM was developed specifically to represent spatial structure and temporal evolution in radar-based precipitation nowcasting (Shi et al., 2015). Applied to this case, such a model could learn the development, movement, splitting and dissipation of convective cells. A third model would estimate forest susceptibility from wind speed and direction, VIL, echo-top height, slope, aspect, valley orientation, topographic exposure, soil moisture, distance from forest edges, tree species and stand structure. An explainability method such as SHAP could show why the model assigns high susceptibility to a particular sector, rather than offering only a black-box score (Lundberg and Lee, 2017). Once a sufficiently large regional archive had been assembled, clustering could group storm trajectories and identify recurrent corridors through valleys, saddles and depressions.

Potential added value and validation requirements. The resulting system could combine forecast storm intensity with local forest susceptibility to estimate the spatial probability of windthrow during the following 30–60 minutes. This would transform the study from reconstruction of a past event into a component of near-real-time forest-risk management. The principal limitation is the rarity of severe events. One storm cannot provide sufficient examples for a robust deep-learning model. Training would require a regional archive covering many Carpathian storms, or transfer learning from comparable mountainous areas followed by local recalibration. The retrospective value of AI in this case is especially clear: it would make operational the future research question already formulated in the original paper.

6. FROM A DETERMINISTIC FLOODPLAIN BOUNDARY TO A PROBABILITY OF INUNDATION

Original research design. The floodplain study linked HEC-HMS hydrological modelling to the one-dimensional MIKE11 hydrodynamic model. The workflow included river axes and cross-sections, roughness coefficients, river-network representation and discharge boundary conditions. A 2 m DEM derived from LiDAR and bathymetric information was used. Cross-sections were generally generated at 100 m intervals, and the article explicitly recognised that modeller subjectivity could intervene decisively in their selection. The comparison between statistical and simulated maximum discharges produced a mean relative error close to 10% (Győri et al., 2016).

AI-assisted redesign. This case should not be approached by replacing hydraulic equations with an unconstrained neural network. It is better suited to theory-guided and hybrid AI (Karpatne et al., 2017; Reichstein et al., 2019). Computer vision applied to high-resolution terrain data could identify the channel axis, banks, bridges, defence structures and abrupt geometrical changes. Cross-sections could then be positioned adaptively: more densely where geometry changes rapidly and less densely in uniform sectors. Orthophotos and LiDAR-derived information could be classified automatically to produce spatially variable estimates of roughness, still subject to hydraulic interpretation and field verification. Bayesian optimisation could calibrate Manning coefficients, boundary conditions and hydrological parameters. After a representative set of MIKE11 simulations had been generated, a surrogate model could learn the relationship among discharge, geometry, roughness and water level. Physics-informed neural networks offer one way of constraining learning with governing differential equations (Raissi et al., 2019), although their suitability would have to be demonstrated rather than assumed.

Potential added value and validation requirements. The most important output would be a probabilistic flood map rather than a single apparently exact boundary. Each location could be associated with a probability of inundation, an expected range of water depths and a decomposition of uncertainty into contributions from discharge, topography, roughness and model structure. The surrogate could permit rapid exploration of bridge obstruction, vegetation growth, channel modification, local defence walls or land-use change. Active learning could indicate the cross-section at which one additional field survey would reduce final uncertainty most efficiently. The original hydraulic model would remain the reference. Any surrogate would require evaluation both inside and outside the parameter combinations used for training, together with tests of mass conservation and hydraulic plausibility. AI would thus transform the map from a deterministic end product into an interactive and uncertainty-aware decision system.

7. FROM MEASURING URBAN IMPERVIOUSNESS TO OPTIMISING URBAN INTERVENTION

Original research design. The urban-runoff study used four Landsat scenes, from 1986, 1993, 2002 and 2014, to identify the expansion of impervious surfaces. Maximum Likelihood, Mahalanobis Distance, Neural Networks and Support Vector Machine classification were compared, and Maximum Likelihood produced the highest accuracy in that experiment. The classified land-cover data were introduced into HEC-HMS. Impervious surfaces increased from 41.7% to 76.1%, accompanied by increases in total runoff and peak flow (Haidu and Ivan, 2016).

AI-assisted redesign. A contemporary redesign would first replace the binary representation of pervious and impervious surfaces with fractional imperviousness. A pixel containing roofs, gardens and pavement would no longer be assigned only to its dominant class. A U-Net-type segmentation model could integrate Landsat, orthophotos, LiDAR and street-network information to distinguish roofs, roads, parking areas, compacted soil, vegetation and semipermeable surfaces (Ronneberger et al., 2015). A multitemporal model would estimate the probable year of conversion for each sector instead of relying on four isolated dates. Hydrological modelling could be extended through radar rainfall, the storm-water network, microtopography, antecedent moisture and multiple observed rainfall events. An AI emulator trained on hydrological simulations would allow thousands of combinations of storms and land-cover configurations to be assessed rapidly. The next methodological step would be multi-objective optimisation of permeable pavements, green roofs, detention basins, rain gardens, urban vegetation and water-storage systems.

Potential added value and validation requirements. The question would change from how increasing imperviousness affected runoff to which urban sectors contribute most strongly to peak discharge and where interventions should be placed to obtain the greatest reduction for a given budget. AI could estimate the marginal contribution of each sector and identify nonlinear thresholds beyond which additional imperviousness generates disproportionate hydrological effects. At the same time, AI should be used to test the robustness of the original conclusion, not merely to reinforce it. The reported correlation was calculated from four temporal configurations subjected to the same rainfall scenario. A richer analysis would require more observation dates, multiple storm events and measured hydrological responses. A complex model must not convert a small-sample association into an unjustified causal statement. The deeper contribution would be the transition from historical impact assessment to spatial optimisation of urban adaptation.

8. FROM ANIMATED TREND SEGMENTATION TO PROBABILISTIC REGIME DETECTION

Original research design. The trend-signal paper proposed an animated method that decomposed a finite time series into sequential sloped trends and displayed the changing signal as new observations accumulated. The method used the Mann–Kendall–Sneyers transformation and a user-defined lag to identify breakpoints. The resulting segments were fitted by least squares, removed from the original series, and the residual was tested iteratively until it approached white noise (Haidu and Magyari-Saska, 2009). In retrospect, this algorithm already contained a principle close to contemporary online learning: the representation of the signal was updated as new data became available.

AI-assisted redesign. The principal extension would be Bayesian Online Changepoint Detection. Instead of selecting a single breakpoint only after the lag criterion is met, the Bayesian procedure continually estimates the probability that a new regime has begun and the posterior distribution of the duration of the current regime (Adams and MacKay, 2007). A switching state-space model, Hidden Markov Model or Gaussian-process model could distinguish changes in slope, mean, variance and autocorrelation. This would help separate persistent trends, abrupt shifts, isolated anomalies and quasiperiodic fluctuations. The lag would no longer depend mainly on manual choice. Simulated series containing known trends, cycles, autocorrelation, noise and breakpoints could be used to select parameters according to detection rate, false-alarm rate, delay and slope error.

Potential added value and validation requirements. The result would include the posterior probability of each breakpoint, uncertainty intervals for the moment of change, probability distributions for segment slopes and a measure of whether the current regime is likely to continue. The original animation could become an operational warning interface. As temperature, precipitation or discharge observations are added, the system could report the probability that a statistically different regime has emerged. Such probabilities would still depend on the assumed observation model, the prior hazard of change and the treatment of long-range dependence. The method would therefore require sensitivity analysis and testing on both synthetic and long instrumental series. This case shows particularly clearly that AI does not appear without intellectual predecessors. Earlier geographical algorithms already anticipated sequential learning, adaptive segmentation and progressive separation of signal from noise; AI would formalise and probabilistically extend that logic.

9. FROM ACCIDENT HOTSPOTS TO DYNAMIC NETWORK RISK

Original research design. The road-accident study analysed 190 accidents registered between 2010 and May 2012. It identified higher frequencies during daytime, rush hours, spring and autumn, and a concentration along the principal urban arteries with a SW–NE orientation. Accident information was collected from local media and integrated with a digitised street network. The article recognised missing information concerning exact causes, vehicle direction and the duration of traffic disruption (Ivan and Haidu, 2012).

AI-assisted redesign. Natural-language processing could extract information automatically from press reports or police records: location, date, time, vehicle type, number of victims, reported cause, weather, manoeuvre and traffic restrictions. Geocoding and entity-matching models could locate events and identify duplicate reports of the same accident. The more important analytical change would be adjustment for exposure. A road with many accidents may be intrinsically dangerous, or it may simply carry much more traffic. Traffic volume, pedestrian flow, vehicle-kilometres and time spent in traffic should therefore be included. The street network could be represented as a graph. Spatio-temporal graph convolutional models were developed to learn connected network structure and temporal dependence simultaneously (Yu et al., 2018). In this context, a network model could use road geometry, intersections, speed limits, lighting, pedestrian crossings, traffic, weather, schools and commercial activities to estimate accident risk for each segment and time interval. Unlike an ordinary Euclidean Kernel surface, it would respect connectivity and direction in the street network.

Potential added value and validation requirements. The output could become a dynamic risk map: the probability of a severe accident on a particular road segment, during a specified time interval and under specified traffic and weather conditions. The model could distinguish risk arising from high exposure from risk associated with geometry, speed, pedestrian interaction, visibility or weather. Causal models could then evaluate a lower speed limit, a redesigned intersection, improved lighting or a new pedestrian crossing by comparing treated segments with similar untreated segments before and after intervention. AI cannot, however, reconstruct with certainty variables that were never observed. Nor can it remove the reporting bias of local media, which are more likely to report severe or unusual accidents. A longer, official and anonymised database would be required, together with explicit safeguards for privacy and fair enforcement. Here AI would transform descriptive hotspot analysis into exposure-adjusted risk estimation and counterfactual evaluation of prevention measures.

10. FROM A BEFORE–AFTER COMPARISON TO A CAUSAL ATMOSPHERIC COUNTERFACTUAL

Original research design. The atmospheric-pollution study examined changes in NO2, HCHO, SO2, O3, CO and aerosol index in ten urban areas of the Grand Est region during the first COVID-19 lockdown. Sentinel-5P observations, ground-monitoring data, meteorological variables, land-cover data and socioeconomic indicators were integrated. An Atmospheric Clearance Index (ACI) was developed, and statistically significant relationships were identified between this index and population density and economic indicators (Kovács and Haidu, 2021). The design compared the lockdown interval in 2020 with the same interval in 2019, examined meteorological homogeneity and constructed ACI by standardising and summing the six pollutant-difference rasters, effectively giving the components equal weight.

AI-assisted redesign. The main AI contribution would concern causal counterfactual construction rather than image classification. A Random Forest or gradient-boosting model could learn the expected concentration of each pollutant from meteorology, seasonality, day of the week and long-term temporal behaviour. The actual 2020 meteorological conditions could then be introduced into this model to estimate the pollution levels expected in the absence of the lockdown. Machine-learning meteorological normalisation has been used precisely to separate weather-related variation from changes associated with emissions or policy (Grange et al., 2018). A Bayesian structural time-series model could construct a second counterfactual from earlier years, comparable regions, mobility, traffic, energy use and industrial activity. Such models estimate an intervention effect by comparing the observed series with a probabilistic counterfactual trajectory (Brodersen et al., 2015). The construction of ACI could also be extended. Rather than assigning equal weights by default, weights could be estimated from a larger European dataset, a latent-variable model, toxicological importance or relationships with health and visibility indicators.

Potential added value and validation requirements. The results could estimate, for each pollutant and city, the change expected from meteorology, the additional change associated with reduced mobility, the contribution of industrial activity, the remaining unexplained component and the uncertainty interval of each estimate. The study could then make statements closer to causal inference while still recognising the observational nature of the evidence. The short TROPOMI archive and the sample of ten cities remain fundamental limitations. A sophisticated AI model cannot create missing pre-2018 satellite observations or a perfect control group. It must express those limitations rather than conceal them. The retrospective redesign would move the analysis from a comparison between two periods towards an explicit estimate of what atmospheric conditions might have occurred without the intervention.

11. CROSS-CASE SYNTHESIS

Although the seven studies concern different phenomena and data structures, their retrospective redesigns converge towards a small number of methodological transformations. **Table 2** summarises the principal transition proposed for each case.

First, AI moves geographical research from rules completely specified in advance towards relationships partly learned from observations. Spectral libraries, thresholds, lags and roughness coefficients remain meaningful, but they can be compared, calibrated or adapted through learning.

Second, AI encourages multimodal integration. Satellite imagery, radar, DEMs, time series, networks, socioeconomic indicators and expert knowledge can be represented within a common analytical workflow. This does not mean that all available data should be used indiscriminately. Their spatial support, temporal resolution, uncertainty and causal status must first be made compatible

Table 2.

Principal methodological transition in the seven retrospective cases.

Case	Original technical core	Principal AI extension	Possible new knowledge / key validation
Forest cover	SAM and GIS change detection	Siamese segmentation and temporal learning	Timing, persistence and type of change; independent temporal ground truth.
Windthrows	NDVI, synoptic maps and radar interpretation	Radar sequence learning and susceptibility modelling	Near-real-time impact probability; multi-event regional archive.
Floodplains	HEC-HMS and MIKE11	Hybrid physical–AI surrogate and Bayesian calibration	Probabilistic inundation; conservation and out-of-sample tests.
Urban runoff	Landsat classification and HEC-HMS	Fractional segmentation and intervention optimisation	Marginal runoff contribution; multiple storms and observed runoff.
Trend signals	MKS segmentation and animated updating	Bayesian online regime detection	Probability and early warning of structural change; benchmark series.
Road accidents	GIS hotspots and temporal frequencies	NLP, graph learning and causal evaluation	Exposure-adjusted dynamic risk; official traffic and accident data.
Air pollution	Satellite differences and composite ACI	Meteorological normalisation and causal counterfactuals	Intervention effect separated from weather; longer series and controls.

.

Third, AI makes it possible to replace a single deterministic result with a probability distribution. A mapped pixel is not merely deforested or unchanged, flooded or unflooded, hazardous or safe. It can be accompanied by a probability, an uncertainty interval and an explanation of the variables that contributed to the result.

Fourth, AI extends research from description towards counterfactual reasoning. It asks not only what occurred, but what would have occurred under another land cover, another policy, another storm or the absence of an intervention.

Fifth, AI links geographical analysis to optimisation. A model can search for the location, timing or combination of interventions that produces the greatest reduction in risk or the largest environmental benefit under explicit constraints.

The deepest benefit is consequently not computational speed. Speed matters, especially where thousands of images or scenarios must be processed, but it is the most superficial advantage. The stronger claim emerging from the seven cases is that AI enlarges the class of geographical questions that can be formulated and made operational.

12. ARTIFICIAL INTELLIGENCE AND CRITICAL GEOGRAPHICAL THINKING

One of the principal objections to AI in education and research is that it may inhibit human critical thinking. This risk should not be dismissed. It can occur when an AI-generated answer is accepted without examination, when references are not verified, when a plausible explanation is confused with evidence, or when a researcher delegates the formulation of the scientific problem itself.

Nevertheless, cognitive inhibition is not an inevitable effect of AI. It depends on the mode of use. In a substitution mode, the human asks for a finished answer and accepts it; here passivity is a genuine risk. In an assistance mode, AI performs repetitive operations while the researcher retains the analytical design. In a collaborative mode, AI generates alternative hypotheses, methods and interpretations that the researcher must compare, challenge, modify or reject. In this third mode, critical thinking is not eliminated. It is required more frequently and at a higher level.

The critical questions also change. The researcher must ask whether the training data are representative; whether spatial autocorrelation or temporal overlap has produced data leakage; whether the chosen scale and resolution are appropriate; whether a model transfers from one region to another; whether a correlation has been interpreted as causation; whether a physical law has been violated; whether uncertainty has been measured; and who may be harmed by a mistaken prediction. Explainability tools can help interrogate a complex model, but they do not transfer responsibility from the researcher to the algorithm (Lundberg and Lee, 2017).

In my own work, I have experienced the opposite of intellectual inhibition. The dialogue that generated this Letter widened the space of methodological alternatives and helped convert an initial intuition into the inverted retrospective paradigm used here. The seven papers were not selected by the AI system, and the link between this experiment, Technical Geography and *antefactum* science did not emerge autonomously from a machine. Those decisions required personal scientific history, editorial responsibility and a conception of where Geography should develop.

A personal experience cannot settle the general educational question, and it should not be presented as universal evidence that AI always increases creativity. It does demonstrate that AI use is compatible with creative geographical reasoning. When critically supervised, it can stimulate questions that might not emerge in the same form through a routine literature review. The relevant educational objective should therefore not be to preserve every manual operation, but to teach students and researchers how to formulate problems, inspect assumptions, verify evidence, compare models and reject unsupported outputs.

Critical thinking should not be measured by the number of calculations performed manually. It should be measured by the quality of the questions, assumptions, tests, interpretations and decisions for which the researcher accepts responsibility.

13. FROM *POSTFACTUM* TO *ANTEFACTUM* GEOGRAPHY

In the 2024 Letter, I described classical Geography as predominantly *postfactum* because it frequently explains phenomena after they have occurred. I argued that Technical Geography should become *antefactum*: predictive, anticipatory and capable of proposing action before damage occurs (Haidu, 2024). The present retrospective experiment shows more concretely how this transition can occur.

Forest research can move from mapping previous losses to forecasting where particular patterns of forest removal may increase hydrological risk. Windthrow research can move from reconstructing storm trajectories to nowcasting the areas that may be affected during the next hour. Floodplain research can move from a fixed hazard boundary to continuously updated probabilities. Urban hydrology can move from documenting increased runoff to selecting the most effective intervention. Trend analysis can move from retrospective segmentation to early detection of a new regime. Accident geography can move from historical hotspots to dynamic segment-level risk. Atmospheric research can move from a before–after comparison to a counterfactual estimate of intervention effects.

This progression may be expressed conceptually as:

Observation → Detection → Explanation → Prediction → Counterfactual Evaluation → Optimisation.

Classical descriptive Geography operated mainly in the first stages.

Technical Geography extended detection, explanation and modelling.

AI can accelerate and connect prediction, counterfactual evaluation and optimisation.

A second conceptual expression may therefore be proposed:

***Antefactum* Technical Geography = geographical reasoning + spatial–temporal data + physical and statistical modelling + learning + uncertainty assessment + decision criteria.**

This is not a mathematical identity. It expresses the components required for a genuinely anticipatory Geography. AI alone cannot create *antefactum* Geography. Prediction without geographical scale, physical meaning, causal interpretation or uncertainty can produce confident but misleading results. Conversely, geographical knowledge without learning and updating may remain unable to process the volume and complexity of contemporary spatial–temporal observations. The *antefactum* potential is generated by their combination.

14. CONCLUSIONS

The retrospective examination of seven earlier studies does not show that their methods were wrong or that their conclusions should be abandoned. Each study addressed a real geographical problem with the data, software and methodological possibilities available at the time.

The exercise demonstrates instead that contemporary AI could have opened additional analytical pathways. In some studies, the principal benefit would probably have been faster classification or calibration. In others, it would have been uncertainty assessment, integration of previously separate datasets, online updating, causal counterfactual construction or optimisation of interventions.

AI's most important contribution is therefore epistemic rather than merely computational. It does not only perform an existing operation more quickly. It may make a new question operational.

Nevertheless, the expected benefit remains conditional. AI requires representative data, independent validation, respect for spatial and temporal dependence, comparison with transparent baselines, physical constraints where appropriate and explicit uncertainty. A complex model should not be accepted because it is complex, and a high predictive score should not replace geographical interpretation.

Human responsibility remains indivisible. AI cannot decide which geographical problem is socially important, whether a mapped pattern is causally meaningful, whether a prediction is sufficiently reliable for public intervention, or whether the costs and benefits of an action are ethically acceptable.

The proper relationship is therefore not Geography versus AI, and not geographer versus machine. It is a relationship in which geographical reasoning defines space, scale, process, causality and validation, while AI extends the capacity to learn, compare, anticipate and explore.

Accordingly, the Editorial Note below sets out the principles that will also be incorporated into the journal's Policy and Guide for Authors and Reviewers.

The future of Geography is not only technical. It is technical, learning-enabled, critically supervised and directed from *postfactum* explanation towards *antefactum* anticipation.

EDITORIAL NOTE: A JOURNAL POLICY FOR RESPONSIBLE AI USE

Geographia Technica permits the responsible use of artificial intelligence and AI-assisted technologies in geographical research and manuscript preparation. AI systems must not be listed as authors. Authors must disclose the name of the tool, the model or version, the purpose of its use and the stages of the research or writing process in which it was employed. Human authors remain fully responsible for the accuracy, originality, references, interpretation and ethical compliance of the submitted work.

When AI is used in data processing, modelling, coding, image generation or modification, or visualisation, its use must be described in sufficient detail in the Methods section to permit scientific evaluation and, where possible, reproducibility. AI-generated or AI-modified visual material must be explicitly identified.

Editors and reviewers must not upload confidential manuscripts, figures, datasets, correspondence or review materials to external AI systems unless the tool has been expressly authorised by the journal and appropriate confidentiality and data-protection safeguards are in place. Any authorised use of AI during peer review must be disclosed to the responsible editor.

This policy supplements the journal's existing adherence to COPE principles and its rules concerning authorship, originality, plagiarism, conflicts of interest and confidentiality. It will be published as a distinct, publicly accessible and versioned section of the journal's Policy and Guide for Authors and Reviewers.

DECLARATION OF AI-ASSISTED WORK

The inverted retrospective paradigm, the selection of the seven case studies, the geographical and epistemological framing, the editorial argument and the final conclusions were conceived, critically evaluated and approved by the author. OpenAI's ChatGPT, operating in Pro mode and powered by GPT-5.6 Sol Pro, was accessed on 23 July 2026 and used in an iterative dialogue to examine the seven published studies; propose AI-enhanced methodological alternatives; compare their potential contributions; organise the cross-case synthesis; and assist with drafting and linguistic refinement. The AI system was not treated as an autonomous source of scientific evidence or as an author. All factual claims, methodological descriptions, references and interpretations were reviewed and verified by the author. AI-generated suggestions were accepted, modified or rejected according to the author's scientific judgment. The author assumes full responsibility for the accuracy, integrity, originality and final content of the published Letter.

REFERENCES

Adams, R.P. & MacKay, D.J.C. (2007) Bayesian Online Changepoint Detection. arXiv:0710.3742. https://doi.org/10.48550/arXiv.0710.3742

Breiman, L. (2001) Random Forests. Machine Learning, 45, 5–32. https://doi.org/10.1023/A:1010933404324

Brodersen, K.H., Gallusser, F., Koehler, J., Remy, N. & Scott, S.L. (2015) Inferring causal impact using Bayesian structural time-series models. The Annals of Applied Statistics, 9(1), 247–274. https://doi.org/10.1214/14-AOAS788

Costea, G., Serradj, A. & Haidu, I. (2012) Forest cartography using Landsat imagery, for studying deforestation over three catchments from Apuseni Mountains, Romania. In: Advances in Remote Sensing, Finite Differences and Information Security, WSEAS Press, 109–114.

Daudt, R.C., Le Saux, B. & Boulch, A. (2018) Fully Convolutional Siamese Networks for Change Detection. Proceedings of the 25th IEEE International Conference on Image Processing, 4063–4067. https://doi.org/10.1109/ICIP.2018.8451652

Furtuna, P., Haidu, I. & Maier, N. (2018) Synoptic processes generating windthrows. A case study in the Apuseni Mountains (Romania). Geographia Technica, 13(2), 52–61. https://doi.org/10.21163/GT_2018.132.04

Grange, S.K., Carslaw, D.C., Lewis, A.C., Boleti, E. & Hueglin, C. (2018) Random forest meteorological normalisation models for Swiss PM10 trend analysis. Atmospheric Chemistry and Physics, 18, 6223–6239. https://doi.org/10.5194/acp-18-6223-2018

Győri, M.-M., Haidu, I. & Humbert, J. (2016) Deriving the floodplain in rural areas for high exceedance probability having limited data source. Environmental Engineering and Management Journal, 15(8), 1879–1887. https://doi.org/10.30638/EEMJ.2016.201

Haidu, I. (2016) What is Technical Geography. Geographia Technica, 11(1), 1–5. https://doi.org/10.21163/GT_2016.111.01

Haidu, I. (2024) Geography's Future is Technical – a letter from the editor. Geographia Technica, 19(2), i–vi. https://doi.org/10.21163/GT_2024.192.23

Haidu, I. & Ivan, K. (2016) The assessment of the impact induced by the increase of impervious areas on surface runoff. Case study: the city of Cluj-Napoca, Romania. Carpathian Journal of Earth and Environmental Sciences, 11(2), 331–337.

Haidu, I. & Magyari-Saska, Z. (2009) Animated Sequential Trend Signal Detection in Finite Samples. Proceedings of the ITI 2009 31st International Conference on Information Technology Interfaces, 249–254. https://doi.org/10.1109/ITI.2009.5196088

Ivan, K. & Haidu, I. (2012) The spatio-temporal distribution of road accidents in Cluj-Napoca. Geographia Technica, 2/2012, 32–38.

Janowicz, K., Gao, S., McKenzie, G., Hu, Y. & Bhaduri, B. (2020) GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. International Journal of Geographical Information Science, 34(4), 625–636. https://doi.org/10.1080/13658816.2019.1684500

Karpatne, A., Atluri, G., Faghmous, J.H., Steinbach, M., Banerjee, A., Ganguly, A., Shekhar, S., Samatova, N. & Kumar, V. (2017) Theory-guided data science: A new paradigm for scientific discovery from data. IEEE Transactions on Knowledge and Data Engineering, 29(10), 2318–2331. https://doi.org/10.1109/TKDE.2017.2720168

Kovács, K.D. & Haidu, I. (2021) Effect of Anti-COVID-19 measures on atmospheric pollutants correlated with the economies of medium-sized cities in 10 urban areas of Grand Est Region, France. Sustainable Cities and Society, 74, 103173. https://doi.org/10.1016/j.scs.2021.103173

Lundberg, S.M. & Lee, S.-I. (2017) A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.

Raissi, M., Perdikaris, P. & Karniadakis, G.E. (2019) Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 378, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N. & Prabhat (2019) Deep learning and process understanding for data-driven Earth system science. Nature, 566, 195–204. https://doi.org/10.1038/s41586-019-0912-1

Ronneberger, O., Fischer, P. & Brox, T. (2015) U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Lecture Notes in Computer Science, 9351, 234–241. https://doi.org/10.1007/978-3-319-24574-4_28

Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K. & Woo, W.-C. (2015) Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. Advances in Neural Information Processing Systems, 28, 802–810.

Yu, B., Yin, H. & Zhu, Z. (2018) Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, 3634–3640. https://doi.org/10.24963/ijcai.2018/505