European Journal for Philosophy of Science (2024) 14:34
https://doi.org/10.1007/s13194-024-00597-2

PAPER IN PHILOSOPHY OF THE NATURAL SCIENCES

Tales of twin cities: what are climate analogues good for?

Giovanni Valente[1] · Hernán Bobadilla[1] · Rawad El Skaf[1] · Francesco Nappo[1]

Received: 2 August 2023 / Accepted: 2 July 2024

Abstract

This article provides an epistemological assessment of climate analogue methods, with specific reference to the use of spatial analogues in the study of the future climate of target locations. Our contention is that, due to formal and conceptual inadequacies of geometrical dissimilarity metrics and the loss of relevant information, especially when reasoning from the physical to the socio-economical level, purported inferences from climate analogues of the spatial kind we consider here prove limited in a number of ways. Indeed, we formulate five outstanding problems concerning the search for best analogues, which we call the problem of non-uniqueness of the source, problem of non-uniqueness of the target, problem of average, problem of non-causal correlations and problem of inferred properties, respectively. In the face of such problems, we then offer two positive recommendations for a fruitful application of this methodology to the assessment of impact, adaptation and vulnerability studies of climate change, especially in the context of what we may prosaically dub "twin cities". Arguably, such recommendations help decision-makers constrain the set of plausible climate analogues by integrating local knowledge relevant to the locations of interest.

Keywords Climate analogues · Climate science · Dissimilarity metrics · Impact · Adaptation and vulnerability assessment

1 Introduction

Analogue methods have a long history of applications in the study of regional future climates (e.g., Glantz, 1988; Darwin et al., 1995; Mendelsohn & Dinar, 1999; Hallegatte, 2007; Veloz et al., 2012). Notwithstanding the empirical limitations they face (Ford et al., 2010; Gutierrez et al., 2019), they are standardly considered as a fruitful addition to regionalized scenario-based computer

✉ Giovanni Valente
giovanni.valente@polimi.it

[1] Department of Mathematics, Politecnico di Milano, via Bonardi 9, Edificio 14 "Nave", Campus Leonardo, Milan 20133, Italy

Published online: 22 July 2024

simulations. In particular, so-called *spatial climate analogues* purport to find locations whose present climate is most similar to the expected future climate of some location of interest. In order to do so, key climate indicators are selected and then, by means of formal metrics of dissimilarity, best analogues for a target region are identified. Such a methodology has been introduced and applied in climate studies at least since the 1980s. For instance, in order to anticipate the effects of global warming for agricultural purposes, Bergthórsson et al. (1988) argued that the current climate of northern Britain could be used as a spatial analogue of the potential future climate of Iceland, so as to draw inferences about the rate of grass growth. Likewise, in more recent developments (e.g., Hallegatte et al., 2007; Pinzon et al., 2021), the methodology has been routinely applied to urban areas with the aim of investigating how future climate change will affect target cities by searching for their best analogues among other, so to speak, twin cities currently experiencing similar climatic conditions to the expected ones. Yet, notwithstanding their significant proliferation in the production of regional climate information, a philosophical analysis of spatial analogue methods is still missing. The question we wish to address in this paper is whether, and to what extent, spatial climate analogues can drive reliable inferences about the future climate of target locations.

The status of climate analogues was systematically reviewed in the 2001 IPCC's third assessment report (TAR), wherein they were discussed as one of three main methods for constructing climate scenarios, together with incremental scenarios and climate model-based scenarios (TAR-WGI: Chap. 13). As such, they were understood as promising a plausible representation of future climate for the purpose of investigating the potential consequences of anthropogenic climate change (ibid., 743). Despite being criticized in their purported function of building scenarios, spatial analogues were nonetheless commended for their value in the assessment of impact, adaptation, and vulnerability (IAV) in climate studies (TAR-WGII, Chap. 18): for, allegedly, they would enable one to transfer experience from existing climatic regions to places where similar climate may be found in the future. To underscore this virtue, subsequent authors of the likes of Hallegatte et al. (2007) even went on to advocate "a heuristic use of 'climate analogues' to circumvent the absence of credible counterfactuals and fully-fledged (sic.) visions of adaptation mechanisms." (p. 2). In spite of the fact that climate analogue techniques effectively treat quantitative data for the relevant climate indicators by means of well-defined dissimilarity metrics, our paper strongly cautions against overly enthusiastic interpretations. In our view, they can be useful as exploratory tools to identify candidate analogues and as means of communicating climate information to policy-makers. However, drawing inferences from them, especially in the assessment of IAV, additionally requires a non-trivial evaluation of the local context of application: indeed, misapplications frequently derive from the attempt to infer directly from best analogues without properly taking into account local knowledge.

In order to arrive at these conclusions, throughout the paper we will raise a series of concerns about the methodology of climate analogues. Specifically, we formulate five outstanding problems, which pose severe limitations on the ability to draw reliable inferences about the target location, at least for the spatial analogues of the kind we consider here. The discussion below is organized as follows. In Section 2, after introducing the concept of spatial analogues, we raise two non-uniqueness problems, one concerning the identification of the source location, namely the alleged best analogue, which is due to the availability of multiple inequivalent dissimilarity metrics (Section 2.1), and the other one concerning the identification of the future climate of the target location, which stems from the uncertainty characterizing the use of climate projections (Section 3.1). The following section is devoted to examining the extent to which, provided that one is able to resolve the previous non-uniqueness problems, fixing a well-defined best analogue could serve as a driver for reliable inferences about the future climate of the target. In doing so, we compare the methodology of climate analogues with the standard approaches to analogical reasoning. That puts us in position to highlight a number of critical issues, taking the form of the problem of average, which is rooted in the use of a geometrical or topological measure of dissimilarity (Section 3.3), as well as of the two related problem of non-causal correlations and problem of inferred properties, which raise further questions about what kind of information may be plausibly transferred from the source to the target (Section 3.2). Finally, in Section 4, we look into the feasibility of spatial analogues in the assessment of impact, adaptation and vulnerability in climate change studies. In particular, in the face of the five problems we formulated, we offer some positive recommendations for a fruitful application of such a methodology, especially in the context of what we may prosaically dub "twin cities". For one, we submit that decision-makers and stakeholders should take into consideration a whole set of climate analogues, rather than fixing on a single one, and then evaluate the plausibility of each of them in light of the specific purposes at stake (Section 3.4). Moreover, they should integrate local knowledge, so as to constrain the set of plausible analogues by further including salient socio-economic as well as cultural factors, which are particularly relevant for cities and other highly populated locations (Section 3.7). We claim that combining together these recommendations with the standard methodology of spatial analogues would prompt an improvement of their use in IAV studies of climate change.

2 Searching for the best climate analogue: dissimilarity metrics and projection uncertainty

Global circulation models are the most widely used tool for projecting future climate scenarios (Chen et al., 2021). Because of their resolution, however, these simulations often prove less helpful as sources of information at a regional or local scale.

To partly compensate for these limitations, a variety of additional methods have been developed in recent years. Climate analogue techniques are one of them.[1] In the literature, two main kinds of analogues are distinguished: *spatial* and *temporal* (Pinzon et al., 2021). Spatial analogues seek to draw inferences about a target location's future climate and vulnerabilities from data about a known source location in the (near) present with similar morphological characteristics and average temperatures. Temporal analogues, instead, look for the deep past of a given region to project its future climate under a variety of physical scenarios. While the latter rely upon paleoclimatic data, typically coming from fossil evidence, whose scarcity and indirect nature may affect the quality of the reconstructions of past climate, the former avail themselves of direct data from a source's present climate and future projected radiative forcing levels, and as such they seem promising for the sake of extracting useful information about the regional impacts of climate change.[2] In what follows, we opt to focus on spatial climate analogues.[3]

The guiding idea of climate analogues of the spatial kind is to look for a source location S whose present climate can be regarded as being the analogue of the projected future climate of a certain target location T. Once S has been identified, typically by means of a measure of dissimilarity taking into account relevant properties, the next question is to what extent information about S can be used to draw reliable inferences about T, both in terms of predictions about physical events (e.g., potentially catastrophic weather and climate conditions) and in terms of impact and effective strategies for adaptation. Arguably, the answer to the second question heavily depends on the answer to the first: in particular, one expects that the inference should be more reliable when the identified source location proves to be a good analogue for the target. In this section, we critically survey the methodology to search for the best climate analogues and then, in the following sections, we address the question of how reliable the purported inferences are as regards the future climate of the target location as well as impacts and adaptation strategies at the socio-economic level.

Clearly, a precondition for the use of spatial analogues is that large amounts of data are available about the climate of the geographical regions of interest. Already on this point, a few preliminary remarks are in order. To begin with, one should note that empirical data are seldom available for exact locations, but rather they are distributed across extended spatial regions. The standard technical procedure to circumvent this issue is to partition the geographical area under study into a large number of grids of equal size, so as to associate the target location with the grid x being centered on it: based on the recorded data, one then searches for another grid

1 Other alternative methods are downscaling and RMCs: cfr. Giorgi (2019) and Tapiador et al. (2020).

2 A potential third kind of climate analogues may be the 'paleoclimatic' ones that are the object of recent studies by Wilson, 2023 and Watkins, 2024, whose purpose is to generate constraints on physical climate models. These climate analogues are not the object of our analysis in what follows. For issues regarding the direct or indirect nature of paleoclimatic data, see Bokulich (2021) and Wilson and Boudinot (2022).

3 For simplicity, henceforth we will use 'climate analogue' (or simply 'analogue') as a shorthand for 'spatial climate analogue'.

$x\prime$ including the best analogue. In addition, let us stress that there is also an issue concerning what indicators, namely the properties about which data are collected, ought to be taken into account for the sake of characterizing the climate (cfr. Werndl, 2016; Frigg et al., 2015; Katzav & Parker, 2018; see also our Section 4 below). More to the point, while there is some consensus on the fact that temperature and precipitations are relevant physical indicators (cf. Kopf et al., 2008), things are less clear when it comes to evaluate the impacts of climate change on a given region: in fact, that seems to be a purpose-dependent matter related to the context of investigation, e.g. impacts on agriculture in the countryside as opposed to flooding in urban cities (cfr. Parker, 2020; Bokulich & Parker, 2021), and there are even aspects of the problem that are hard to quantify, such as cultural attitudes and lifestyle. Be that as it may, once the grid structure is fixed and the relevant indicators are selected, data are collected during regular intervals of time (typically on a monthly basis) for a sufficiently long period (typically of the order of 30 years). It should be stressed that direct observational data are only gathered for the present climate of the candidate source locations, whereas for the data representing the future climate of the target location are supposed to be projected values of the indicators, which are computed by means of computer simulations of some chosen climate model and a chosen forcing level.

Based upon the collected data, a quantitative method is then introduced in order to evaluate how similar the numerical values for the climates of the source and the target are for each indicator. This methodology is rooted in the so-called ‘geometrical’ or ‘topological’ view of similarity, whereby one adopts a mathematically well-defined measure to calculate the distance between two systems, such as T and S, formally conceived as points in a metrical space. As we will argue in the present section, the prospect of finding the best climate analogue for the target location is limited by a number of factors. For one, there are different dissimilarity metrics adopted in the literature for the sake of determining whether a certain location be analogous to the target. All such metrics have rather distinct statistical features and, as it turns out, they typically disagree on their conclusions. Thus, in the absence of any empirical way to establish which measure is correct, it becomes rather problematic to identify the alleged best analogue. There thus arises a *problem of non-uniqueness for the source* S. Furthermore, since the expected future climate of the target location is posited on the basis of climate projections, it necessarily depends on the choice of the models and the forcing levels that are run in the simulations. There thus arises also a *problem of non-uniqueness for the target* T. That is aggravated by the fact that the projected outcomes are affected by deep uncertainty, both within a single model and across the different models. In what follows, we discuss in greater detail these two distinct non-uniqueness problems.

2.1 Non-uniqueness of the source

The non-uniqueness of the measure of dissimilarity adopted to quantify the extent to which the climate of some candidate location can be regarded as analogous to the climate of the target location is a well-known issue. For instance, Grenier et al. (2013)

reviewed and contrasted six of the main dissimilarity metrics employed in the literature (that is, the *Standardized Euclidean Distance*, the *Kolmogorov-Smirnov Statistic*, the *Nearest-Neighbor Distance*, the *Zech-Aslan Energy Statistic*, the *Friedman-Rafsky Runs Statistic*, the *Kullback-Leibler Divergence*) and showed explicitly how their conclusions about the alleged best analogues of a given target diverge in a concrete case. Specifically they look for candidate source locations within the entire North America to be analogues of the city of Montreal. Moreover, the authors introduced a number of a priori and a posteriori criteria for metrics comparison and evaluated whether or not each metric fares well with them. We will not recall all such criteria, the more so because most of them hinge upon computational considerations; nor do we need to survey all the dissimilarity metrics listed by Grenier et al. Yet, it is worth focusing on two of them, namely the Nearest-Neighbor Distance and the Standardized Euclidean Distance, that are particularly relevant for our conceptual analysis. Specifically, we will discuss how they perform under the a priori criterion of "continuous discrimination", which reflects the ability of a metric to determine an unequivocal ranking of the candidate locations for best analogue. Such a provision is motivated by the need to unambiguously identify the most suitable climate analogue of the target: for, if two distinct candidate locations have the same degrees of (dis-)similarity from the target despite having different climates from each other, then one would find oneself without a definite basis for the purported analogical inference. As a matter of fact, if a metric has discrete range, it can hardly satisfy the criterion of continuous discrimination.

In order to spell out the formulas for the dissimilarity metrics of interest, we now introduce some basic elements of the formalism. Let the climate at the target location and the climate at the candidate source location be represented by the vector distributions $\psi_T = \{X_1, \dots, X_n\}$ and $\psi_S = \{Y_1, \dots, Y_n\}$, respectively, where the number n of points in each distribution corresponds to the number of years (or months if one wants to take into account more frequent data) in which the quantities X_i and Y_j are measured. Such quantities belong to a real-valued space whose dimension d coincides with the total number of climate indicators. Thus, for (the grid of) each location one considers an ensemble of $n \times d$ collected data, which describe relevant aspects of the climate during a given period of time. Allegedly, the climates of the target and the candidate location may then be viewed as analogues if the measure of dissimilarity computed between the respective distributions ψ_T and ψ_S has a very low value, meaning that they are not too dissimilar. Of course, this fact depends just on what metric one adopts.

Let us begin with the Nearest-Neighbor Distance. The first step is to find for each point in the two distributions its nearest neighbour within the pooled distribution (i.e., the ensemble of all the elements of the two individual distributions taken together). That is supposed to be the closest point with respect to the Standardized Euclidean distance between points, which has the following form:

$$\mathrm{SED}\left(\mathrm{X_i}, \mathrm{Y_j}\right) = \sqrt{\sum_{\mathrm{k}=1}^{\mathrm{d}} \frac{\left[\mathrm{X_i(k)} - \mathrm{Y_j(k)}\right]^2}{\sigma_{\mathrm{T}}(\mathrm{k})\sigma_{\mathrm{S}}(\mathrm{k})}}$$

Accordingly, for any indicator indexed by k, one first computes the square of the difference between the points X_i and Y_j, divided by the product $\sigma_T(k)\sigma_S(k)$ of the standard deviations of ψ_T and ψ_S with respect to the given indicator (which assures that the various k's are given the comparable weight); and then one sums the results over all indicators and finally takes the square root of the sum. So, for each point X_i in ψ_T one looks for another point in the pool distribution for which $SED\left(X_i, \bullet\right)$ has the smallest value, and symmetrically for each point $pointY_j in \psi_S$ one looks for the point that minimizes the distance $SED\left(\bullet, Y_j\right)$. In order to establish the degrees of dissimilarity one counts the overall number N_{NN} of points whose nearest neighbour belongs to the same original distribution. This procedure is meant to capture how dissimilar ψ_T and ψ_S are from each other in the sense that, if the closest neighbour of a point in ψ_T belongs again to the target distribution or if the closest neighbour of a point in ψ_S belongs again to the candidate distribution, then such points contribute to N_{NN}; otherwise, they do not contribute to N_{NN}. The Nearest-Neighbor Distance is then defined as the ratio of the latter quantity and the number of points in the pool distribution:

$$\mathrm{D_{NN}} = \frac{\mathrm{N_{NN}}}{2\mathrm{n}}$$

As such, it quantifies the relative number of points for which the closest point is in its original distribution, and hence a low value tells us how distant, or mutually dissimilar, the target and the candidate distribution are. Note that the lower bound of the thus-defined dissimilarity metric is 0, which obtains when each point in a distribution has its closest neighbour in the other distributions, and hence ψ_T and ψ_S are, so to speak, maximally similar; whereas its upper bound is 1, meaning that the two distributions are maximally dissimilar, in that none of their points have its closest neighbour in the other distribution. The range of D_{NN} is necessarily discrete: specifically, it comprises $2n + 1$ values, each one separated by an interval $1/2n$. A direct consequence of this fact is that the Nearest-Neighbor Distance cannot satisfy Grenier et al.'s (2013) criterion of continuous discrimination, since it is possible that the climate distributions of different candidate locations have the same value of D_{NN} when being computed with respect to the target distribution.

The next dissimilarity metric has the virtue of overcoming such a limitation. In this case, instead of focusing on the distance between pairs of points, one calculates the distance between the two distributions taken on average. The Standardized Euclidean Distance between the averages of the distributions has the following form:

$$\mathrm{D_{SED}}\left(\psi_\mathrm{T}, \psi_\mathrm{S}\right) = \sqrt{\sum_{\mathrm{k}=1}^{\mathrm{d}} \left\{ \frac{\left[\psi_\mathrm{T}\right](\mathrm{k}) - \left[\psi_\mathrm{S}\right](\mathrm{k})}{\sigma_\mathrm{T}(\mathrm{k})} \right\}^2}$$

where the operator $[\psi](k)$ denotes the average of all n points in a given distribution ψ with respect to the dimension k. In order to avoid possible confusion, let us reiterate that, even though they are both types of Standardized Euclidean Distance, the previously introduced SED, which serves to define the metric D_{NN}, is relative to points, whereas the metric D_{SED} introduced in the above formula is relative to averages of distributions. Accordingly, one computes the square of the difference between the average of ψ_T and ψ_S divided by the standard deviation of the target distribution ψ_T, and then one sums the results over all indicators $k = 1, \ldots, d$ before taking the square root of the overall sum. The thus-defined dissimilarity metric yields an index of the extent to which the two distributions diverge on average for each indicator. It is bounded from below by 0, which obtains when for all indicators the average values are equal, meaning that ψ_T and ψ_S are maximally similar. The range of D_{SED} then grows continuously as the averages of the two distributions become more and more dissimilar for each dimension. This entails that the Standardized Euclidean Distance, contrary to the discrete Nearest-Neighbor Distance, fares well with the criterion of continuous discrimination. Indeed, one can unambiguously rank distinct candidate distributions on the basis of the value of $D_{SED}(\psi_T, \bullet)$ calculated with respect to the target distribution. Note, however, that the standardization procedure contemplated in the above formula is enacted with respect to the target distribution ψ_T and not even with respect to ψ_S, as in the denominator there appears only the standard deviation $\sigma_T(k)$. That implies that, differently from the usual Euclidean distance, the present version of the Standardized Euclidean Distance is non-symmetrical, in the sense that in general one has $D_{SED}(\psi_T, \psi_S) \neq D_{SED}(\psi_S, \psi_T)$. Hence, it has the consequence that, even though a certain location may be identified as the best analogue of the target, the latter does not need to be the best analogue of the former.

As we already pointed out, there are various inequivalent measures of dissimilarity available, which yield rather different results when being applied to concrete cases. The example of the city of Montreal, evaluated under the six dissimilarity metrics Grenier at al. (2013) consider, illustrates this fact: for instance, the analysis of the authors shows that, when projections for the future climate of Montreal are obtained with SIM-01 for the years 2041–2070, most of the locations ranked within the analogues in the range 200–500 are found in Nova Scotia according to D_{SED}, while that is not the case according to D_{NN} and the other metrics. The inequivalence between the different measures of dissimilarity is what leads one to the problem of non-uniqueness of the source. Let us note here that the Euclidean Distance is the most basic example of a metric used to quantify how close two systems are with respect to a given set of known properties. In fact, even the other dissimilarity metrics for climate analogues involve some kind of statistical average of the differences between the values of T and S across all relevant climate indicators. As such, it suffices to focus on the Standardized Euclidean Distance defined above for the sake of illustrating the 'geometrical' or 'topological' view of similarity, whose philosophical significance will be discussed in detail in Section 3.

2.2 Non-uniqueness of the target

Next, we wish to draw attention to another problem affecting the methodology of spatial climate analogues, which introduces deep uncertainty concerning the proper identification of the target (rather than of the source). It is that the future climate of the target location is determined indirectly on the basis of model-dependent climate projections, and as such it is not uniquely fixed. Indeed, while the points in the distribution ψ_S of the candidate source location are informed by actual observational data, the points in the target distribution ψ_T correspond to expected outcomes of model simulations under certain forcing level. This gives rise to what we call the problem of non-uniqueness of the target T. In particular, the underdetermination of the future climate of the location under study can be traced back to three distinct, albeit interrelated, factors: that is, the consideration of different scenarios under which one projects the present state of the location; the use of multiple, typically inequivalent, models predicting the evolution of the climate in the course of time; and the fact that there are various layers of uncertainty surrounding the whole computational process. Let us elaborate on this non-uniqueness problem in the remaining part of the section.

Climate projections are conceived as predictions conditional to a certain level of radiative forcing, which are formulated on the basis of emissions scenarios. Levels of radiative forcing are indeed the key input for Global Circulation Models (GCM) to project future climates at a global scale and, by downscaling, future climates of regional targets as well. What matters for our purposes here is to note that there are a variety of emissions scenarios that are then translated into specific forcing levels, moving from the original SRES (Special Report on Emissions Scenarios) for climate change, which are no longer in use, to the recent generation of scenarios, which mainly combine two ingredients, namely the RCPs (Radiative Forcing Pathways) and the SSPs (Shared Socioeconomic Pathways). As it happens, in the literature on climate analogues, there is no standard choice of emissions scenario, nor of the forcing level to adopt: for instance, Kopf et al. (2008) looked at future targets under SRES A2, which generates an RCP close to 8.5, while Pinzon et al. (2021) used RCP 6.0. However, the expected data informing the points in the target distribution ψ_T are highly sensitive to the choice of the level of radiative forcing. In fact, different levels project different future climates for any given target location. Thus, deciding on the use of a specific RCP (i.e. level of radiative forcing) turns out to be crucial for the sake of making predictions about the future climate of a spatial region of interest. In addition, as it will become more evident in Section 4, when it comes to evaluate the impacts of climate change and to design adaptation strategies one should also take into account the socio-economic context of the target location, which can be particularly complex for highly populated urban areas, like a city. For this purpose, it seems convenient to combine RCPs with SSPs in order to formulate relevant projections. Nevertheless, there are several different SSPs available that are consistent with any given RCP, thereby projecting the target into rather different futures[4]. Due to the proliferation of possible scenarios, the problem of non-uniqueness of the target presents itself.

[4] In fact, Shared Socioeconomic Pathways comprehend five alternative characterizations of plausible societal futures, positing different scenarios for the socio-political as well as institutional organization of the world: namely, SSP1 (sustainability), SSP2 (middle of the road), SSP3 (regional rivalry), SSP4 (inequality), and SSP5 (rapid growth).

Such a problem is not just a consequence of considering multiple scenarios. Indeed, even by keeping a fixed level of radiative forcing, if one employs different Global Circulation Models, one obtains a variety of projected climates, which would in turn have distinct spatial analogues. For example, Hallegatte et al. (2007) show that, in order to simulate the future climate of Paris, the resulting analogue is Bordeaux under the ARPEGE model whereas it is Cordoba under the HadRM3H model. The non-uniqueness of the target thus arises also due to the availability of multiple models under which climate simulations are performed. What is more, the model-dependence of projections further complicates the interpretation of the target distribution ψ_T. In principle, we see two possible ways of understanding the meaning of the distribution points: either they refer to a single model m, so that one should more properly denote them by the labels $X_1[m], \dots, X_n[m]$; or they are intended with respect to an ensemble of models, so that the projected data in the set are computed as average values $\overline{X}_1, \dots, \overline{X}_n$ over all models. Yet, under both understandings one faces thorny difficulties. In the first case, the target distribution turns out to be highly sensitive to what specific RCP level is built into the chosen model, which has drawbacks for the decisions to be taken by policy makers. Instead, in the second case, the establishment of the target distribution and hence the search for its best analogue acquire an additional statistical component, over and above the use of any specific dissimilarity metric averaging across climate indicators.

Besides the multiplicity of scenarios as well as climate models, the underdetermination of the future of the target location is further aggravated by the fact that the process of building projections is affected by deep uncertainty of various kinds. To begin with, it is customary to distinguish projections from standard (unconditional) predictions, in that projections are sensitive to initial conditions uncertainty. The latter stems from our lack of knowledge about what exactly one should take as the present state of the model from which simulations are launched (see Werndl, 2019 for a critical analysis of this claim). Moreover, since the projected outcomes are relative to the specific choice of a climate model, they inherit the uncertainties of the model itself. For one, there is structural uncertainty about the form that modelling equations should take on, as well as parametric uncertainty about the numerical values ascribed to the salient parameters (cf. Parker, 2018a, 2018b). Consequently, one cannot predict or project how the climate will precisely evolve in the course of time. In order to control such uncertainties, one technique consists in considering ensembles of climate models so as to determine average expected outcomes. Yet, even though it is true that robustness analysis can aid to extract reliable information from ensembles (cfr. Lloyd, 2010; O'Loughlin, 2021; Parker, 2018a, 2018b; Winsberg, 2018), this entails that the projections about the future of the target only have a purely statistical character, thereby raising a number of epistemic questions concerning whether ensemble results should be regarded as providing probabilistic facts and, if so, how we should go about to infer probabilities from them[5] (cf. Bishop

[5] Other questions about the interpretation of climate projections can be posed, which are still debated in the philosophical literature: e.g., do they provide a "non-discountable envelope", in the sense of a lower bound of future climate changes (Stainforth et al., 2007)? or should they be regarded as real possibilities to be taken seriously (Katzav, 2014)?

& Abramowitz, 2013). In addition to the just mentioned structural and parametric uncertainty, as Bradley et al. (2017) emphasized, it is also unclear how to apply the outcomes of global models about the Earth's climate to a local scale. This form of uncertainty is particularly relevant for the issue at stake here, since one needs to determine the projected values of climate indicators for a local region containing the target location. Lastly, on top of the physics of the climate, when investigating the future of populated geographical areas, one ought to take into consideration socio-economic factors as well, which are somewhat elusive and thus introduce further layers of projections uncertainty (we will return to this in Section 4).

All in all, the availability of different projected futures for a given location depending on the choice of a relevant RCP and circulation model is the root of the problem of non-uniqueness of the target, which in itself is independent from the problem of non-uniqueness of the source we stated in the previous sub-section. Although there are attempts to cope with uncertainty in the context of the search for best spatial analogues (e.g., Pinzon et al., 2021), it remains an outstanding issue to predict how the climate of any location of interest will evolve in the incoming decades, thereby making it more complicated to determine what other location may presently have comparable climatic conditions, irrespective of which measure of dissimilarity one opts for. We will discuss a possible way to assuage both uniqueness problems in Section 4, when dealing with the applicability of the methodology of climate analogues to IAV studies.

3 What kind of analogical reasoning?

Provided that, in the face of the two above-mentioned non-uniqueness problems, one has managed to identify the best climate analogue of a given target location, the next question to address is whether, and to what extent, one can draw reliable inferences. In order to answer such an outstanding question, it is worth reviewing the basic concepts of analogical reasoning, so as to contrast them with the methodology of climate analogues of the spatial kind that is employed in scientific practice. As we show in the present section, this puts us in a position to raise additional problems concerning the use of climate analogues, over and above the problems of non-uniqueness of the source and of the target.

Analogical inference is a distinctive form of reasoning in science as in everyday life, whereby one draws conclusions about a target system based on its similarities with some source system. According to the standard schematization, the source system S is supposed to possess a set of properties p_k's, with $k = 1, \dots, d$, which are similar to the corresponding properties p_k^*'s possessed by the target T; then, if S possesses some extra property p_{d+1}, one should infer that T also possesses the corresponding property p_{d+1}^*. The outstanding question is whether, and when, such conclusions are plausible. To answer such a question in full generality, the required degrees of similarity should be specified for each property: that is, how similar should they be to count as positive analogies? Moreover, besides having d properties in common, the two systems typically will feature some known negative analogies, i.e., properties that are not similar, as well as neutral analogies, i.e., properties for which we do not know whether there is any similarity (see Bartha, 2009 for a systematic treatment of analogical reasoning). So, one may ask: how is the strength of the conclusion about the additional $d+1$ property affected by the negative and neutral analogies? And to what extent can the positive analogies, especially if there are a large number of them, be sufficient to drive the inference?

A venerable approach to analogical reasoning, which traces back to Hesse (1963), attempts to provide a framework that, while renouncing to answering questions about analogical inference in full generality, aims to make agreement and disagreement about such inferences in science intelligible to both practitioners and outsiders. From this viewpoint, Hesse argues that the evaluation of a scientific analogy occurs along two distinct dimensions at once. First, evaluation occurs at the level of how similarities and dissimilarities are expressed: claims to the effect that two systems are similar or dissimilar in some respects require genuine respects of similarity, rather than artificially introduced ones aiming solely to inflate an otherwise implausible analogy. Secondly, an analogical argument is regarded as strong only if there is a causal connection among some of the properties p_k's that can also be expected among the corresponding properties p_k^*'s, which should be relevant for the additional $d+1$ property in the conclusion. The advantage of this approach is to yield an informative account of how analogical reasoning is assessed in practice while deferring to the scientific parties engaged in investigation the difficult questions about just how much similarity is needed for a given conclusion to be regarded as plausible.

An influential view in contemporary epistemology extends this basic framework to all "surrogative reasoning" in science, i.e., to any inferential circumstance in which some system S is taken as a 'surrogate' (though not necessarily an analogue) of a target. Weisberg (2012, 2013), in particular, has argued that the reliability of surrogative inferences about the target is grounded by the fact that a scientific model provides an adequate representation of the target in virtue of their being similar in relevant respects. Underlying Weisberg's approach is, first of all, a notion of representation that implies resemblance between target and surrogate model. Although Weisberg goes on to defend a specific account of resemblance (the so-called 'contrast approach'), his notion presupposes, at a minimum, the existence of a one-to-one correspondence between the properties of surrogate and target. Moreover, Weisberg's approach presupposes a notion of relevance for similarities that is plausibly spelled out in causal terms: as Weisberg writes, representational capacity is related to ensuring that "a logical consequence of a model is mirrored as a causal consequence in the world" (2013:169). However, the contrast approach has been criticized by several authors working on scientific modelling (see for instance Frigg & Nguyen, 2020), and we ourselves contend that a geometric approach to similarity appears to be a more suitable description for the methodology of spatial analogues. In fact, the dissimilarity metrics adopted in the climate literature are just meant to quantify the topological distance between source and target computed on average among a set of properties. Thus, as we argue below, they do not lend weight to the one-to-one correspondence between properties; nor do they retain any causal significance, in that they merely establish statistical correlations. It follows that, differently from the contrast approach to similarity, within the geometrical approach the source has rather weak representational power over the target, and hence the ensuing surrogative inferences appear to be much less grounded.

In order to see this, let us now compare the above-described scheme of analogical reasoning with the use of spatial climate analogues. In the latter, the relevant properties to compare target and source are given by the selected indicators k's. A first difference with standard analogical inference concerns the number of indicators one takes into account. In analogical reasoning, the list of potentially relevant characteristics is often open-ended and its definition is left to the reasonable disagreement between scientific parties. While one expects an analogical inference to become stronger as one finds more and more positive analogies, very few indicators are employed in the search for best climate analogues. In particular, it is commonly believed (based partly on the influence of Holdridge's (1947) 'life zone' system) that three physical indicators, of which at least one related to temperature and another related to moisture, are sufficient in order to adequately characterize the climate of a certain location. For instance, Kopf et al. (2008) consider solely annual data about aridity, heating degree and cooling degree days of certain regions of the world in order to estimate those regions' adequacy as spatial climate analogues for target urban areas.

Another difference one should keep in mind is that, while in standard analogical reasoning one takes into account the one-to-one similarities between individual properties of the target and source, in dealing with climate analogues the degrees of similarity are computed by appealing to a metric function that averages across all indicators. That is, instead of evaluating each indicator individually and retain those for which the respective values for the source and the target are numerically close, one applies a well-defined measure to determine an overall dissimilarity score, where all relevant indicators are assigned equal weight. The lack of an explicit one-to-one comparison between the relevant properties of T and S leaves it open that, even for the best analogue, there are some climate indicators for which the difference with the target location are extremely large. In particular, the Euclidean Distance is the basic example of how a region S could be selected as best analogue for a target T in spite of a significant negative analogy – so long as it is compensated by other positive analogies. In fact, virtually all measures of dissimilarity adopted in the literature prescribe averaging out the values of selected indicators, in line with the geometrical approach to surrogative reasoning.

3.1 The problem of average

The above described methodological setting for the choice of best climate analogues raises further epistemological concerns, over and above the non-uniqueness problems discussed in Section 2. Granted that inferences from climate analogues present some differences from ordinary analogical inferences, one might still ask: to what extent could scoring high on a given formal dissimilarity metric possibly drive any inference at all from the source's climate to a target? As we have seen, even when we abandon the intuitive notion of similarity and consider the more general category of surrogative reasoning, it seems plausible to require that there be some expected match between the causal connections across the properties that hold in the source system and those that hold among the corresponding properties in the target system. Indeed, the very notion of *relevant* similarities presupposes a story about how the predicted property p_{k+1} is robustly connected, in the surrogate, to the properties p_k's that a surrogate shares with the target (Weisberg, 2013). By averaging each region's score across each indicator, instead, and by considering only small sets of indicators, spatial analogues can be insensitive both to potentially critical differences between source and target with respect to properties included as key indicators and to potentially critical factors not countenanced among the key indicators in the model. This gives rise to what we can call the *problem of average*. As an illustration, consider Kopf et al.'s (2008) study of best climate analogues for twelve European cities under the assumption of radiative forcing of the A2 SRES scenario. The authors identify best analogues for each city by considering aggregate average score over the three key indicators of aridity, heating degree days and cooling degree days. One concern about this method is that differences in one of the key indicators, e.g., aridity, may not be well-compensated by similarities across the other two indicators, even though the difference is partially blunted when averaging. For instance, a value for aridity beyond a certain threshold may well have dramatic effects on a given urban area that are not easily comparable to those on another region regarded as similar on average, in the sense of having a very low value for the chosen dissimilarity metric, but whose aridity level is still below the threshold. The problem of average can thus lead to serious consequences.

There is no easy solution to such a problem. In their review of the dissimilarity metrics for climate analogues, Grenier et al. (2013) invoke an evaluation criterion that they dub "balance of 1D departures", whereby differences between the target and its putative analogue ought to be (roughly) the same across all one-dimensional indicators taken individually. This condition is intended to rule out cases in which the candidate location with the lowest dissimilarity value has a high difference with respect to one indicator that is nonetheless counterbalanced by low differences for the other indicators. Indeed, in such a case the resulting overall best analogue would be a bad analogue for at least one property. The requirement that all 1D departures are roughly the same guarantees that the source and the target are comparatively similar in all relevant respects. While this is certainly useful, it should be noted that such a criterion can only be applied a posteriori: that is, it is only after one identifies the best analogue on the basis of a chosen dissimilarity metric that one can evaluate whether it exhibits balance across all indicators. Even so, though, based on data relative to a specific regional study, Grenier et al. (2013) provide a comparison of how various extant dissimilarity metrics fare with this a posteriori criterion, finding significant variability in the results. What is more, balance across all indicators is not sufficient in itself to guarantee that the resulting best analogue is good enough, since it may happen that the magnitude of 1D departures is very large. Yet, if one imposes an acceptability bound on the magnitude, then there is no guarantee that a best analogue is found.

Another possible solution to the problem of average can be elaborated along the lines of the approach put forward by Hallegatte et al. (2007). It consists in choosing a preferred indicator and then imposing acceptability bounds on the others. More to the point, the authors consider a 30-year monthly means for both temperature and precipitations for the period 1960–1990 and compute projections for the future period 2070–2100. They then employ three distinct distance metrics: dT encoding the mean absolute difference between the temperature monthly means; d_P^A denoting the relative difference between the annual mean precipitations, which accounts for the total water availability; and d_P^M denoting the mean relative difference between monthly mean precipitations, which is instead relevant for the impacts on infrastructures and lifestyles. Hallegatte et al. claim that dT ought to be regarded as privileged since the source and target are supposed to be similar at least with respect to their average temperature. However, they also stress that this cannot be the unique indicator to take into account, for if monthly mean temperature is considered individually it tends to select locations that are spatially very close to the target or to pick analogues that differ significantly in terms of precipitations. For this reason, Hallegatte et al. introduce tolerance margins using the other two metrics, specifically they required 15% for d_P^A and 30% for d_P^M. Their proposed criterion for the best analogue is thus to identify the spatial grid that minimizes dT while remaining within the tolerance margins for d_P^A and d_P^M, as long as dT does not itself exceed the 1 K degree bound. Let us note that, in this approach, the magnitude of the tolerance margins is merely stipulated, thereby introducing some degrees of arbitrariness in the search for best analogues. Thus, in this approach it is not just the choice of the relevant indicators but also that of their respective acceptability bounds that is submitted to expert judgement. Furthermore, a possible drawback of adopting this selective approach is that there may well be no acceptable analogue at all for a given region of interest. For instance, Hallegatte et al. point out that for the city of Geneva the ARPEGE-Climat model predicts that, when the precipitations constraints $d_P^A < 15\%$ and $d_P^M < 30\%$ are fulfilled, the lower value for the mean absolute difference between the temperature monthly means is 3.08 K, which clearly exceeds the acceptability margin for dT.

All in all, the problem of average makes the search for best analogues quite elusive, even when one fixes a specific distance metric to capture the degrees of dissimilarity between the source and the target. For, on the one hand, the insensitivity of any such metric to the one-to-one correspondence of specific indicators can lead one to identify as the best climate analogue a candidate location that has undesirable properties, despite showing the lowest score of dissimilarity with the target. On the other hand, the available solutions of the problem that seek to add further constraints to the condition of minimal dissimilarity run into the risk of failing to identify a candidate location that can adequately serve as the best analogue. What is more, it is reasonable to expect that good candidates for spatial analogue, especially when looking for twin cities, should be compared with the target not only with respect to physical indicators of the climate, but also to quantitative indicators connected with socio-economical drivers. We conclude this section by pointing out that the problem of average is even closely related to other two problems concerning the choice of relevant indicators over which one computes the statistical average prescribed a specific measure of dissimilarity. Let us spell out these further problems below.

3.2 The problem of non-causal correlations and the problem of inferred properties

The next problems we formulate cast doubts on the ability of the best analogue, however it is defined according to the adopted dissimilarity metric, to drive reliable inferences from its present climate to the future climate of the target location.

To begin with, let us stress that there arise some issues dealing with causality. For one, the causal mechanism that has led to the current climate in the source location need not to be the same as the mechanism for the expected climate change in the target. In fact, the future climate is simulated by means of long-term projections, say of the order of 30 years, which, besides being subject to deep uncertainty, are also a function of the particular level of radiative forcing that is supposed to act on Earth for the next three decades. However, this level does not coincide with the forcing level that has effectively acted on Earth for the past three decades, which instead contributed to change the climate of the candidate location into what it is nowadays. When this happens, the evolutions of the target's and the source's climates do not share the same causal driver. Even worse, as a consequence of the problem of average, the degrees of dissimilarity between the target location and its best analogue appear as void of causal significance, inasmuch as the geometrical distance metrics in use only establish statistical correlations across physical indicators. Indeed, averaging over the numerical differences between the relevant climate indicators one ends up neglecting possible causal connections among them. By contrast, causal correlations are regarded as crucial in order to draw reliable inferences within standard schemes for analogical reasoning. Let us elaborate on this point.

Recall that according to accounts such as the one proposed by Hesse (1963), an analogical inference is enforced by the existence of a causal connection between some of the properties p_k's of the source system which is reflected onto the corresponding properties p_k^*'s of the target system: if such properties are among the positive analogies between S and T, and the property p_{k+1} possessed by S is also connected to them, then one is licensed to infer that T possesses (with some plausibility) the corresponding property p_{k+1}^*. Instead, in the case of climate analogues, the indicators k over which one computes a measure of dissimilarity between the source and the target are not really expected to be correlated with each other. For example, temperature means, or even the number of cooling or heating days, is supposedly independent from the amount of precipitations, or at least one does not explicitly assume any causal connection between these indicators. Given that the measure of dissimilarity just indicates the presence of statistical correlations between the indicators k taken on average, it remains unclear whether, and in what sense, a low score in the chosen dissimilarity metric could drive any inference from the present climate of the best analogue to the future climate of the target location. This is what we call the *problem of non-causal correlations*. Granted, that is a puzzling issue just for those accounts of analogical inferences that demand that the drivers should have a causal underpinning; but, in any case, it is still revealing of the purely statistical character of the conclusions made by reasoning on the basis of dissimilarity metrics that average across mutually independent factors.

There is yet another issue that depends on the choice of the relevant set of indicators, which is relevant for the question whether the use of climate analogues can license one to transfer additional properties, or even just gather any useful information, to the target location. As before, the comparison with the standard accounts of analogical reasoning gives some insight. In the latter, the purpose of an analogical inference is to support the claim that the target system T possesses some additional property p^*_{k+1}, granted that the analogue source system S possesses a similar property p_{k+1}. By contrast, in the procedure to search for climate analogues, one does not include any additional indicator $k + 1$ corresponding to the property to be inferred. As a result, it becomes unclear how to specify the properties that one could plausibly transfer from the present climate of the source location to the future climate of the target location. More to the point, such properties should have some kind of connection with the relevant indicators k's. However, if the indicators are taken to be all physical properties that are supposed to characterize the climate of a certain region, then there is a great deal of uncertainty about how to transfer other physical, ecological or even socio-economical properties typical of the present state of one particular city to the future state of another city. And of course the issue becomes even more problematic when it comes to transfer cultural habits. That is what we refer to as the *problem of inferred properties*. It becomes a pending matter inasmuch as, besides elaborating formal methodologies to identify the best spatial analogues, the literature seldom explicates what exact information one should extract from them that can be conveyed to the target locations.

In the last analysis, the problem of non-causal correlations and the problem of inferred properties pose severe limitations on the ability to make analogical inferences about the effects of climate change on target locations on the basis of their best climate analogues. Arguably, the properties that one would like to transfer depend on the purpose of the inference. It is widely recognized by climate scientists that different choices of key indicators tend to work best for different purposes (see Parker, 2020 for a recent contribution on the topic). In fact, the process of selecting a suitable set of relevant indicators seems to be typically guided by mere experience, thereby rendering the issue how to draw reliable inferences hard, if not impractical, to settle a priori. That is even clearer in light of the fact that a great deal of recent applications of the methodology of climate analogues has to do just with the possible impacts of climate change, especially in vulnerability and adaptation studies. In the next section, we move on to survey these applications.

3.3 Climate analogues as tools for assessing impact, adaptation, and vulnerability

As we pointed out in the Introduction, climate analogues were listed in the TAR-WGI among the different types of scenarios of future climate in use. On this point, it is worth stressing that the report reached a rather critical conclusion about the status of spatial analogues as climate scenarios:

> The approach is severely restricted by the frequent lack of correspondence between other important features (both climatic and non-climatic) of a study

> region and its spatial analogue (Arnell et al., 1990). Thus, spatial analogues are seldom applied as scenarios, per se. (748)

The systematic analysis of the methodology of climate analogues we delivered in previous sections enforces and extends the negative conclusion expressed by the TAR, at least for the spatial analogues of the kind we have considered. More precisely, the five outstanding problems we have formulated above affect the use of any alleged best analogue as a reliable representation of the future climate of the target location. The two non-uniqueness problems call into doubt the possibility of properly identifying the source and the target, whereas the problem of average and the closely related problems of non-causal correlations and inferred properties challenge the ability to draw reliable inferences from the source to the target. In light of such limitations, there arises the following question, which inspires the title of the present paper: What, if anything, are climate analogues good for?

To answer this question with a positive approach, at least as regards spatial analogues, it is insightful to see how the TAR-WGI continues its critical conclusion cited above:

> Rather, [spatial analogues] are valuable for *validating* the extrapolation of *impact models* by providing information on the *response of systems* to climatic conditions falling outside the range currently experienced at a study location. (748, our emphasis).

Indeed, it is mostly as a tool for assessing impact, vulnerability and adaptation (IAV) that spatial analogues have been used. This is also emphasised in TAR-WGII, in particular in Chap. 18 dedicated to adaptative measures to climate change in the context of sustainability and equity: "Knowledge of the processes by which individuals or communities actually *adapt* to changes in conditions over time comes largely from *analog* and other empirical analyses" (887-8, our emphasis). Thus, the hope is that, for the sake of planning an effective course of action in response to future climate change at a target location, one could try to extract useful information from the current state of affairs at some other similar location. Granted, the role of climate analogues in the assessment of IAV has been progressively reduced in the latest IPCC report (in fact, they are not even cited in AR6). Yet, it is also true that research on spatial climate analogues has continued to develop, especially to anticipate the effects of climate change on rural areas for agricultural purposes as well as on urban areas for the sake of city development strategies (e.g., Fitzpatrick & Dunn, 2019, Hancock et al., 2017).

In this section, we argue that research on climate analogues should continue to be developed for the sake of advancing the assessment of IAV and beyond. Our specific approach is to appraise the spatial analogues methodology, but with adequate amendments that are meant to help one cope with its outstanding problems. To this aim, in the next subsections we make two positive recommendations. First, in the face of the underdetermination of the best analogue for the target location, decision-makers and stakeholders should consider a whole set of analogues and evaluate how plausible each of them may be with respect to the specific purpose at stake. Second, the search for climate analogues should be guided by local knowledge, in such a way to integrate physical and socio-economic factors, as well as behavioural and cultural elements, which are particularly relevant in the case of highly populated areas. We submit that combining these recommendations would improve the use of climate analogues of the spatial kind in IAV studies: while the first opens up the space of solutions to include different climatically plausible analogues, the second narrows down this set by using local knowledge in plausibility assessment, including socio-economic factors, thereby strengthening analogical reasoning.

3.4 First recommendation: evaluating a set of plausible climate analogues

At the core of our reflection on how climate analogues could be effectively applied to concrete situations is the fact that IAV studies deal with various sorts of climate risks. Stakeholders and decision-makers typically assess the impact of risks relative to their local contexts. In particular, they establish strategic priorities based on their cities' vulnerabilities and engineer risk-specific adaptation plans. In this sense, we suggest that the ultimate unit of analysis in the assessment of IAV is what one may call *city risks*, intended as the particular risks that decision-makers are concerned with when designing adaptation measures for their own cities. Of course, this does not mean that IAV should not aim at a comprehensive planning that may address the totality of risks an urban area is exposed to. Rather, it only means that decisions and adaptation strategies are relative to the risks at hand: indeed, it is not uncommon for city planners to opt for prioritization in risk management or even focus exclusively on critical risks, given financial or temporal constraints. Confronting themselves with a presently observable analogue of what their city may become in the incoming decades due to climate change can thus put them in a position to better assess potential risks. The virtuous usage of analogical tools is explicitly underscored by authors such as Hallegatte et al. (2007): as they put it, "the distance between the projected city and its current analogue provides a synthetic indicator of the importance of the adaptation to be operated and a basis for more in-depth analyses of the risks involved in various transition paths toward a fully re-adapted city." (p.6).

Yet, this valuable prospect for advancing IAV management appears to be doomed, at least *prima facie*, by the underdetermination of the best spatial analogue for the location of interest, which derives from both the problem of non-uniqueness of the target and the problem of non-uniqueness of the source. For one, the plurality of scenarios under which climate simulations are performed on the basis of different global circulation models, aggravated by the various layers of projection uncertainty that we described in Section 3.1, would not yield to decision-makers a univocal picture of what the future climate of their city may look like. What is more, as explained in Section 2.1, even if they were to select a specific projection, for instance by taking a risk-averse attitude whereby one adheres to the most extreme scenario of severe climate change (i.e. the highest level of radiative forcing fed into the most pessimistic GCM), different dissimilarity metrics may as well point them to different best analogues. As a result, stakeholders would not be able to recognize the sought-after twin city to serve as a reference in their IAV assessment for policy-making.

However, here we wish to assuage worries of non-uniqueness. As a matter of fact, climate experts nowadays tend to regard the availability of multiple different scenarios as a richness rather than a weakness. Faced with uncertainty about the future, it is indeed advisable to take into consideration a larger suite of alternatives, instead of banking on a single scenario. More to the point, the increasingly employed SSP-RCP framework features a variety of possible futures for the world, incorporating also socio-economic elements related to the most significant drivers of climate-related risks in diverse sectors such as health, food, water, and infrastructure, which taking collectively spans an expanded range of uncertainty concerning both mitigation and adaptation measures. Such a pluralistic perspective is particularly appealing in the context of IAV assessment, which involves the ponderation of views coming from a wide array of stakeholders that usually have disparate interests, beliefs and values. In fact, no single city may prove to be the best analogue along different dimensions of risk assessment. For example, the current climate of Cape Town in South Africa is shown by Pinzon et al. (2021) to be the most similar in terms of the combination of surface-average temperature and precipitation to the projected climate of Santiago in Chile, based on computer simulations of MRI-AGCM3.2 H model under the SRES A1B. But if Cape Town and Santiago do not share similar vulnerabilities, analogical reasoning for IAV purposes is rather limited in scope, thereby failing to enact reliable inferences. And for that matter it may as well happen that alternative candidate cities having good, albeit less optimal, dissimilarity values, e.g. some other regions located in selected areas of Australia and South America according to the metric adopted by Pinzon et al., can serve as more useful sources of information for Santiago. In light of this, we suggest that if decision-makers are presented with a manageable set of locations with analogue climates to the one projected for their target city, then they may look for those candidates that exhibit similar vulnerabilities, possibly with the aid of experts on different risks, so as to inspect suitable adaptation strategies.

So, our first recommendation for a virtuous application of the spatial analogue methodology to IAV studies is that, rather than searching for a single best climate analogue, decision-makers should take into account a whole set of possible analogues, as they arise from multiple projections of the future of the target location in compliance with the condition of low dissimilarity scores according to well-defined metrics. Of course, not all candidates may be reliable, or even just useful for the purposes at stake, and hence some caveats are in order. For one, there are practical restrictions, in that by relaxing the requirement to select a unique source the set of alternatives may expand to an intractable extent. To prevent this from happening, a pragmatic upper-limit cutoff on the number and kinds of locations to be taken as references can be placed on the basis of the temporal and financial limitations faced in the process of policy-making. Furthermore, and perhaps more importantly, the set of possible analogues should be constrained by plausibility considerations, in the sense that despite having low dissimilarity scores some putative source locations may fail to be similar to the target in crucial respects. That requires one to deal with the problem of average and the critical fact that focusing on the physical indicators for the climate overlooks the important of socio-economical factors that are relevant to risk assessment for IAV studies. We take this issue up in the rest of the paper.

Before doing so, let us mention that in a recent contribution to the philosophical literature Wilson (2023) addresses a uniqueness problem concerning paleoclimatic analogues. In that context, the problem is that the observed change of modern climate has a unique dynamics that cannot be straightforwardly recovered by comparison with the climate of past periods. Wilson's proposed resolution consists in implementing Currie's (2018) strategy of "exquisite corpse", by which one can combine different partial paleoclimatic analogues so as to provide an adequate reconstruction of the historical evolution of the climate. It should be noticed, though, that the underdetermination of the best spatial analogue we are concerned with in our paper is importantly different from the apparent problem affecting paleoclimatic studies: for one, our focus is on local geographical target regions rather than on the evolution of the climate at a global scale; moreover, the non-uniqueness of the source stems from the availability of distinct dissimilarity metrics, which can give inequivalent results. That said, we actually share the spirit of Wilson's proposal. Indeed, the candidate source locations ought to be regarded as partial analogues of the projected future of the target inasmuch as, even though they are supposed to be similar to it on average, they all exhibit some specific disanalogies too. Based on this recognition, our recommendation of evaluating a manageable set of climate analogues with variations in terms of risks, instead of identifying an alleged best analogue, provides a more comprehensive understanding of the potential effects of climate change on the location of interest.

On this point, what remains to explicate is how the partial climate analogues worth including in the set could meet standards of plausibility. The problem of average reveals that, if the chosen topological metric prescribes one to average over all selected indicators, one can hardly justify the provision of low dissimilarity score as a driver for plausible analogical inferences. As we explained in Section 3.3, a possible solution goes along the lines of the criterion put forward by Hallegatte et al's (2007), according to which one first compares the target with potential source locations with respect to what one takes to be the primary indicator, and then one proceeds to set acceptability bounds for the other indicators. One drawback, though, is that, while tolerance margins can be established in such a way to avoid extreme values that may reach dangerous tipping points, one should still provide reasons as to why one particular indicator is granted a privileged role. Arguably, the required justification would vary from context to context, and thus on the specific purposes of policy-making can provide guidance on what factors one should prioritize. As an illustration, stakeholders could be more concerned with rising temperatures in urban areas where air conditioning plans need to be evaluated; or they could be more concerned with water scarcity in more rural areas characterized by extensive agricultural activity. So, the extent to which one may plausibly infer indirect information about the target location from climate analogues determined by diverse dissimilarity metrics is dependent on the relevant interests of IAV studies.

Another related issue is that the standard methodology for spatial analogues we have reviewed concentrates mainly on physical similarities between the projected climate of the target location and the present climate of the source. Regrettably, in doing so, it disregards salient socio-economic factors, as well as behavioural and cultural ones. However, the impacts of climate change depend not only on physical factors but also on the future vulnerabilities of societal structures that will experience climatic changes in the future and on their adaptative capacities. To be sure, as Ford et al. (2010) note, it is typically presupposed that analogue cities are "also expected to share similarities in terms of socio-economic-political organization." But this expectation stands as a vague desideratum and fails to be explicitly articulated in the methodologies based on dissimilarity metrics that compute physical indicators. Adequate sources should instead be chosen relative to the risks attended in the target location. And, while average temperature and precipitation might be more or less reliable proxies for the climate, they are not so for all types of risks. Therefore, given that similar climates in cities with different vulnerabilities lead to different city risks, spatial analogues with low dissimilarity scores do not need to be good climate risk analogues. We submit that, for the sake of constraining the set of candidate sources so as to offer plausible analogues to decision-makers for specific risks, one further needs to integrate local knowledge of physical and socio-economical contexts of interest. That leads us to our second recommendation.

3.5 Second recommendation: integrating local knowledge and socio-economic factors

Let us recall that the related problem of non-causal correlations and problem of inferred properties that we discussed in Section 3.2 highlight the need to provide a firm ground for analogical inferences from the present state of the source location to the projected future of the target by taking into account more refined information about the relevant local contexts. On this point, we can use Kopf et al.'s (2008) work as telling example. In order to justify their choice of physical indicators for the purpose of analysing the climate impact on cities and urban life, these authors assume that the key indicators of heating and cooling degree days correlate with energy demand. Surely, it is reasonable to expect that the need for household temperature regulation, which plausibly changes on the basis of the outside weather, has a causal connection with energy consumption. One could thereby hope to infer information useful to estimate the future prices of goods and services in the target cities. However, the extent to which Kopf et al.'s causal assumption is supposed to hold varies on the basis of the local context of each city. Indeed, energy demand is a function of socio-economic factors, such as cost and availability of energy, as well as of other contextual factors, such as architectural structures, building materials, quality of heating systems, and even cultural habits of the urban population. Integrating such factors into analogical reasoning clearly requires resorting to local knowledge.

The term "local knowledge" broadly refers to various forms of epistemic goods that are produced by human communities as a result of their social practices in direct interactions with their surroundings (cf. Klenk et al., 2017; Naess, 2013; Dekens, 2007; Canagarajah, 2002). It typically involves all sorts of information, from the physical climate and its observed consequences on the environment to socio-economic and cultural elements that are relevant to IAV studies, thereby showing how local communities adapt to their changing environments (including successes and failures) and the experienced impacts of the changing climate on their patterns of behaviour. As such, it has already been used in the literature to advance climate adaption (see, e.g., Makondo & Tomas, 2018; Hiwasaki et al., 2015; Aswani et al., 2015; Nakashima et al., 2012). In fact, decision-makers may be interested in addressing extreme climate events that are not exclusively influenced by temperature and precipitation. For example, while floods are certainly affected by precipitation indicators, they also depend on local variables, such as topography and mechanical properties of the soil. Likewise, forest fires are directly proportional to increases in temperature, but their occurrence depends on a host of other local variables, from natural to anthropogenic. In this sense, decision-makers need to be aware of how local variables may interact with the projected climate of their target cities, thus producing specific hazards. This is where local knowledge can be usefully put at work.

The idea behind our second recommendation for the improvement of the spatial analogues methodology is that by looking further into the local contexts, one can find wide-ranging and more detailed information, often of high quality, which enables a more systematic comparison between the target and a candidate analogue. Indeed, local knowledge of socio-economic properties and cultural habits can reveal novel (dis)similar properties between them, over and above the physical information encoded in dissimilarity metrics. More to the point, we see two possible ways in which additional information made available by local knowledge can effectively foster the reliability of analogical inferences from source locations to the target. As we argue, while the first possibility faces some hindrances, the second one can be interpreted as giving a plausibility criterion to constrain the set of climate analogues that decision-makers can evaluate for their assessment of IAV studies of climate change. By following the first proposal, one may encode additional information about local communities into quantitative properties that can serve as indicators to fed into a dissimilarity metric, along with average temperature and precipitations. Accordingly, spatial analogues will be determined by averaging directly over both physical and socio-economic indicators via the standard methodology. However, it should be noted that, although there are indeed properties that can be adequately quantified, e.g. energy costs, there other salient properties, especially those relative to population's habits, that do not readily lend themselves to a formal treatment, and as such they could not be factored in. Another limitation of this proposal is that by restricting oneself to the standard dissimilarity metrics one would run again into the problem of average.

The second proposal is instead more promising. It prescribes that relevant information acquired through local knowledge be used just after the candidates for climate analogues are selected by having low dissimilarity values of the adopted measure of physical dissimilarity, so as to further deepen the comparison with the projected future of the target location. The criterion for enhancing plausibility that we advocate thus aims to combine the methodology of spatial analogues with standard analogical reasoning we discussed in Section 2. Accordingly, if an additional property is shared between source and target, the purported inferences would seem more reliable, modulo the already mentioned limitations of analogical reasoning; whereas, if the property is not shared, that would provide reason to discard the candidate location as a proper spatial analogue for the target location, in spite of its low dissimilarity score. In order to enforce such a selective criterion, let us invoke a further condition for the critical properties coming from local knowledge to be considered: that is, that they should not only be relevant but also robust in the sense of remaining stable under extended periods of sustained interactions with the surroundings. That enables one to take into account even more elusive properties like cultural habits, provided that the local communities maintain them in the course of time. Arguably, the more numerous are the relevant and robust properties that are shared, and the more confident one should feel in drawing conclusions about the target on the basis of the analogue source. As the proposed criterion is applied to each one of the candidate locations with low dissimilarity scores, one can restrict the set of potential analogues to those that appear more plausible for the purpose of IAV assessment. In this manner, our two recommendations for concrete applications of the methodology of spatial analogues are fruitfully employed in conjunction.

Before concluding, we would like to stress a further virtue of resorting to local knowledge, as well as a caveat. On the positive side, it can enhance the actionability of ensuing adaptation strategies. Actionability is indeed a central desideratum in climate research, but one that remains largely unattained. Local knowledge advances it in terms of affording procedural legitimacy and meaningfulness. As for procedural legitimacy, the inclusion of views and concerns of stakeholders as part of the development of spatial analogues for adaptation to climate change makes the purported inferences not merely a technocratic tool for finding solutions but also a significantly more democratic process (at least when the input is incorporated in the right way). The perception of the population as being part of democratic and just procedures advances the prospects for effective implementation of adaptation strategies (see, e.g., Brink et al., 2023; Baker & Constant, 2020; Fischer, 2000). As for meaningfulness, by being identified together with local stakeholders, climate analogues integrate key elements of the local experiences, knowledge and values to which the communities at stake can relate. This does not only improve communication and understanding of the ensuing adaptation strategies, but also increases risk awareness about aspects of climate change that are meaningful to the population (cf. Shepherd & Lloyd, 2021; Shepherd et al., 2018; Jasanoff, 2010).

Concerning the caveat, recall that there is a temporal asymmetry between the source and the target in that, while for the former one looks at the present state, for the latter one ought to take into consideration its projected future. As a consequence, local knowledge of a potential analogue can be acquired directly. Instead, present observations of the local climate and socioeconomic structure of the target location can only be used to build statistical expectations for its future state. As recently noted by Pirani et al. (2024), this poses the "limitation… that local, bottom-up studies are often based on understanding current exposure and vulnerability, and then inferring that risk will be higher in the future because climate hazards will increase. This approach doesn't account for possible future changes in vulnerability" (Pirani et al., 2024). So, one would wonder whether climate projections really offer any help in assessing the future vulnerability of target cities. The recent employment of Shared Socioeconomic Pathways seem to provide a promising avenue, in that they incorporate political and socio-economical elements in scenario-building processes. Yet, they describe projected futures at the global scale, and hence they can hardly give insight onto local areas. On a more positive note, it should be reported that global SSPs have also served as a basis for more refined or "extended" SSPs, which operate at the regional, national, and sub-national scales (see, e.g., Rohat et al., 2019). Such extended SSPs may actually provide useful information about the future target in terms of projected population growth, their location in large cities, or in coastal zones, as well as other socioeconomic factors that are relevant to the assessment of IAV.

4 Conclusion

In this paper, we have developed a critical analysis of class of spatial climate analogues, which adopt topological dissimilarity metrics to identify source locations that purport to mimic the projected future climate of given target locations. We formulated five outstanding problems affecting such a methodology. In the face of the limited scope of the analogical inferences one could draw about the target, we offered two positive recommendations for a fruitful application of spatial analogues in the assessment of impact, adaptation and vulnerability studies of climate change, which consist in presenting decision-makers with a whole set of candidate analogues subject to plausibility constraints informed by a recourse to local knowledge.

Needless to say, fostering the reliability of a particular analogical inference is not tantamount to rendering its conclusions certain. As with all forms of non-deductive reasoning, analogical inferences can afford valuable conclusions, but with inherent limitations. Given these limitations, the epistemic goods attained through analogy must be used cautiously, especially when it comes to socioeconomic decisions with cascading effects. If an adaptation policy works in the source domain and relevant similarities are established with the target, this does not imply that the policy will work in the target location, since the analogical inference can be defeated by a series of other climatic and non-climatic factors. This issue is raised, albeit more generally for any science-based decision, by Nancy Cartwright in her 2012 paper "Will this policy work for you", where she argues that, for policies in a source location to have chances to work also in a target location, they must have the right capacities in coordination with specific features of the target. Our own analysis of the methodology of spatial analogues should be intended as contribution in this direction in the context of policy-making for climate change.

Funding Open access funding provided by Politecnico di Milano within the CRUI-CARE Agreement.

References

Arnell, N., Brown, R., & Reynard, N. (1990). *Impact of climatic variability and change on river flow regimes in the UK*. Institute of Hydrology Reports.

Aswani, S., Vaccaro, I., Abernethy, K., Albert, S., & de Pablo, J. F. L. (2015). Can perceptions of environmental and climate change in island communities assist in adaptation planning locally? *Environmental Management,56*(6), 1487–1501. https://doi.org/10.1007/s00267-015-0572-3

Baker, S., & Constant, N. L. (2020). Epistemic justice and the integration of local ecological knowledge for marine conservation: Lessons from the Seychelles. *Marine Policy,117*, 103921.

Bartha, P. (2009). *By parallel reasoning*. Oxford University Press.

Bergthórsson, P., Björnsson, H., Dýrmundsson, Ó., Gudmundsson, B., Helgadóttir, A., & Jónmundsson, J. V. (1988). The effects of climatic variations on agriculture in Iceland. In M. Parry, T. Carter, & N. Konijn (Eds.), *The impact of climatic variations on agriculture*. Springer.

Bishop, C., & Abramowitz, G. (2013). Climate model dependence and the replicate earth paradigm. *Climate Dynamics,41*(3–4), 885–900.

Bokulich, A. (2021). Using models to correct data: Paleodiversity and the fossil record. *Synthese,198*(Suppl 24), 5919–5940.

Bokulich, A., & Parker, W. (2021). Data models, representation and adequacy-for-purpose. *European Journal for Philosophy of Science,11*, 1–26.

Bradley, R., Helgelson, C., & Hill, B. (2017). Climate change assessments: Confidence, probability, and decision. *Philosophy of Science,84*(3), 500–522.

Brink, E., Falla, A. M. V., & Boyd, E. (2023). Weapons of the vulnerable? A review of popular resistance to climate adaptation. *Global Environmental Change,80*, 102656.

Canagarajah, S. (2002). Reconstructing local knowledge. *Journal of Language Identity and Education,1*(4), 243–259.

Cartwright, N. (2012). Will this policy work for you? Predicting effectiveness better: How philosophy helps. *Philosophy of Science,79*, 973–989.

Chen, D., Rojas, M., Samset, B. H., Cobb, K., Diongue Niang, A., Edwards, P., Emori, S., Faria, S. H., Hawkins, E., Hope, P., Huybrechts, P., Meinshausen, M., Mustafa, S. K., Plattner, G. K., & Tréguier, A. M. (2021). Framing, context, and methods. In V. Masson-Delmotte, P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, & B. Zhou (Eds.), *Climate Change 2021: The physical science basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (pp. 147–286). Cambridge University Press.

Currie, A. (2018). *Rock, bone, and ruin: An optimist's guide to the historical sciences*. The MIT Press.

Darwin, R., Tsigas, M., Lewandrowski, J., & Raneses, A. (1995). *World agriculture and climate change: economic adaptations.* Report to the US Department of Agriculture, AER-703.

Dekens, J. (2007). Local knowledge for disaster preparedness: A literature review. *International Center for Integrated Mountain Development*.

Fischer, F. (2000). *Citizens, experts, and the environment: The politics of local knowledge*. Duke University Press.

Fitzpatrick, M. C., & Dunn, R. R. (2019). Contemporary climatic analogs for 540 North American urban areas in the late 21st century. *Nature Communications,10*, 614.

Ford, K., Keskitalo, E., Smith, T., Pearce, T., Berrang-Ford, L., Duerden, F., & Smit, B. (2010). Case study and analogue methodologies in climate change vulnerability research. *Wiley Interdisciplinary Reviews: Climate Change,1*(3), 474–392.

Frigg, R., & Nguyen, J. (2020). *Modelling nature: An opinionated introduction to scientific representation*. Springer International Publishing.

Frigg, R., Thompson, E., & Werndl, C. (2015). Philosophy of climate science part I: Observing climate change. *Philosophy Compass,10*(12), 953–964.

Giorgi, F. (2019). Thirty years of regional climate modeling: Where are we and where are we going next? *Journal of Geophysical Research: Atmospheres,124*(11), 5696–5723.

Glantz, M. (Ed.). (1988). *Societal responses to regional climatic change: Forecasting by analogy*. Routledge.

Grenier, P., Parent, A., Huard, D., Anctil, F., & Chaumont, D. (2013). An assessment of six dissimilarity metrics for climate analogues. *Journal of Applied Meteorology and Climatology,52*(4), 733–752.

Gutierrez, J., Maraun, D., Widmann, M., Huth, R., Hertig, E., Benestad, R., Roessler, O., Wibig, J., Wilcke, R., Kotlarski, S., San Martín, D., Herrera, S., Bedia, J., Casanueva, A., Manzanas, R., Iturbide, M., Vrac, M., Dubrovsky, M., Ribalaygua, R., … Pagé, C. (2019). An intercomparison of a large ensemble of statistical downscaling methods over Europe: Results from the VALUE perfect predictor cross-validation experiment. *International Journal of Climatology,39*(9), 3750–3785.

Hallegatte, S., Hourcade, J., & Ambrosi, P. (2007). Using climate analogues for assessing climate economic impacts in urban areas. *Climatic Change,82*, 47–60.

Hancock, G., Verdon-Kidd, D., & Lowry, B. (2017). Soil erosion predictions from a landscape evolution model – an assessment of a post-mining landform using spatial climate change analogues. *Sciences of the Total Environment,601*, 109–121.

Hesse, M. (1963). *Models and analogies in science*. University of Notre Dame.

Hiwasaki, L., Luna, E., Syamsidik, & Marçal, J. A. (2015). Local and indigenous knowledge on climate-related hazards of coastal and small island communities in Southeast Asia. *Climatic Change,128*, 35–56.

Holdridge, L. (1947). Determination of life plant formation from simple climatic data. *Science,105*(2727), 367–368.

Jasanoff, S. (2010). A new climate for society. *Theory Culture & Society,27*(2–3), 233–253.

Katzav, J. (2014). The epistemology of climate models and some of its implications for climate science and the philosophy of science. *Studies in History and Philosophy of Modern Physics,46*, 228–2382.

Katzav, J., & Parker, W. S. (2018). Issues in the theoretical foundations of climate science. *Studies in History and Philosophy of Modern Physics,63*, 141–149.

Klenk, N., Fiume, A., Meehan, K., & Gibbes, C. (2017). Local knowledge in climate adaptation research: Moving knowledge frameworks from extraction to co-production. *Wiley Interdisciplinary Reviews: Climate Change,8*(5), e475.

Kopf, S., Ha-Duong, M., & Hallegatte, S. (2008). Using maps of city analogues to display and interpret climate change scenarios and their uncertainty. *Environment and Urbanization,29*(1), 195–216.

Lloyd, E. A. (2010). Confirmation and robustness of climate models. *Philosophy of Science,77*(5), 971–984.

Makondo, C. C., & Thomas, D. S. (2018). Climate change adaptation: Linking indigenous knowledge with western science for effective adaptation. *Environmental Science & Policy,88*, 83–91.

Mendelsohn, R., & Dinar, A. (1999). Climate change, agriculture and developing countries: Does adaptation matter? *World Bank Research Observer,14*, 277–293.

Naess, L. O. (2013). The role of local knowledge in adaptation to climate change. *Wiley Interdisciplinary Reviews: Climate Change,4*(2), 99–106.

Nakashima, D. J., Galloway McLean, K., Thulstrup, H. D., Ramos Castillo, A., & Rubis, J. T. (2012). *Weathering uncertainty: Traditional knowledge for climate change assessment and adaptation*. UNESCO, and UNU.

O'Loughlin, R. (2021). Robustness reasoning in climate model comparisons. *Studies in History and Philosophy of Science Part A,85*, 34–43.

Parker, W. (2018a). Climate Science, in Zalta, E. (Ed.). *The stanford encyclopedia of philosophy* (Summer 2018 Edition).

Parker, W. S. (2018b). The significance of robust climate projections. In E. A. Lloyd, & E. Winsberg (Eds.), *Climate modelling* (pp. 273–296). Springer International Publishing.

Parker, W. S. (2020). Model evaluation: An adequacy-for-purpose view. *Philosophy of Science,87*(3), 457–477.

Pinzon, R., Nakaegawa, T., Hibino, K., & Takayabu, I. (2021). A climate analogue approach to understanding the future climates of six western South American capital cities. *Atmòsfera,34*(3), 255–266.

Pirani, A., Fuglestvedt, J., Byers, E., O'Neill, B., Riahi, K., Lee, J., Marotzke, J., Rose, S., Schaeffer, R., & Tebaldi, C. (2024). Scenarios in IPCC assessments: Lessons from AR6 and opportunities for AR7. *npj Climate Action*. 3(1). https://doi.org/10.1038/s44168-023-00082-1

Rohat, G., Flacke, J., Dosio, A., Dao, H., & van Marseveen, M. (2019). Projections of human exposure to dangerous heat in African cities under multiple socioeconomic and climate scenarios. *Earth's Future,7*(5), 528–546.

Shepherd, T. G., & Lloyd, E. A. (2021). Meaningful climate science. *Climatic Change,169*(1–2), 17.

Shepherd, T. G., Boyd, E., Calel, R. A., Chapman, S. C., Dessai, S., Dima-West, I. M., Fowler, H. F., James, R., Maraun, D., Martius, O., Senior, C. A., Sobel, A., Stainforth, D. A., Tett, S. F. B., Trenberth, K. E., van den Hurk, B. J. J. M., Watkins, N., Wilby, R. L., & Zenghelis, D. A. (2018). Storylines: an alternative approach to representing uncertainty in physical aspects of climate change. *Climatic Change,151*, 555–571.

Stainforth, A., Allen, M., Tredger, E., & Smith, L. (2007). Confidence, uncertainty and decision-support relevant in climate predictions. *Philosophical Transactions of the Royal Society A,365*, 2145–2161.

Tapiador, F. J., Navarro, A., Moreno, R., Sánchez, J. L., & García-Ortega, E. (2020). Regional climate models: 30 years of dynamical downscaling. *Atmospheric Research,235*, 104785.

Veloz, S., Williams, J., Blois, J., He, F., Otto-Bliesner, B., & Liu, Z. (2012). No-analog climates and shifting realized niches during the late quaternary: Implications for 21st-century predictions by species distribution models. *Global Change Biology,5*, 1698–1713.

Watkins, A. (2024). Using paleoclimatic analogues to inform climate projections. *Perspectives on Science,1–45*, 1. https://doi.org/10.1162/posc_a_00622

Weisberg, M. (2012). Getting serious about similarity. *Philosophy of Science,79*(5), 785–794.

Weisberg, M. (2013). *Simulation and similarity: Using models to understand the world*. Oxford University Press.

Werndl, C. (2016). On defining climate and climate change. *The British Journal for the Philosophy of Science,67*(2), 337–364.

Werndl, C. (2019). Initial-condition dependence and initial-condition uncertainty in climate science. *British Journal for the Philosophy of Science,70*(4), 953–976.

Wilson, J. (2023). Paleoclimate analogues and the threshold problem. *Synthese,202*(1), 17.

Wilson, J., & Boudinot, F. G. (2022). Proxy measurement in paleoclimatology. *European Journal for Philosophy of Science,12*(1), 14.

Winsberg, E. B. (2018). What does robustness teach us in climate science: A re-appraisal. *Synthese,198*(Suppl 21), 5099–5122.