



SDG interlinkage networks: Analysis, robustness, sensitivities, and hierarchies

J.H.P. Dawes

Centre for Networks and Collective Behaviour and Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, UK



ARTICLE INFO

Article history:
Accepted 3 September 2021

Keywords:
Sustainable Development Goals
Network science
Mathematical modelling
Trade-offs
Co-benefits
Synergies

ABSTRACT

A growing literature considers the Sustainable Development Goals (SDGs) as a interlinked network, connected by co-benefits and trade-offs between pairs of SDGs. Such network descriptions naturally prompt important questions concerning the emergence and identification of system-level features.

This paper develops mathematical techniques to address, quantitatively, the extent to which these interlinkage networks point to the likelihood of greater progress on some SDGs than on others, the sensitivity of the networks to the addition of new links (or the strengthening or weakening of existing ones), and the existence of implicit hierarchies within Agenda 2030.

The methods we discuss are applicable to any directed network but we interpret them here in the context of three interlinkage matrices produced from expert analysis and literature reviews. We use these as three specific examples to discuss the quantitative results that reveal similarities and differences between these networks, as well as to comment on the mathematical techniques themselves. In broad terms, our findings confirm those from other sources, such as the Sustainable Development Solutions Network: for example, that globally SDGs 12–15 are most at risk.

Perhaps of greater value is that analysis of the interlinkage networks is able to illuminate the underlying structural issues that lead to these systemic conclusions, such as the extent to which, at the whole-system level, the structure of SDG interlinkages favours some SDGs over others. The sensitivity analyses also suggest ways to quantify possible improvements to an SDG interlinkage network, since the sensitivity analyses are able to identify the modifications of the network that would best improve outcomes across the whole of Agenda 2030. This therefore indicates possibilities for informing policy-making, since the interlinkage networks themselves are implicitly descriptions of the overlaps, co-benefits and trade-offs that are anticipated to be likely to arise from a set of existing or proposed future policy actions.

© 2021 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The United Nations' Sustainable Development Goals (SDGs), agreed by the UN General Assembly in September 2015 (United Nations, 2015) are recognised as providing a blueprint for a prosperous, sustainable, and equitable future for humanity and as a result set out an ambitious and wide-ranging agenda across all fields of human endeavour. The SDGs are purposefully described as a single 'integrated and indivisible' agenda (United Nations, 2015, Declaration, paragraph 6), with universal applicability, rather than as a collection of independent ambitions across the three dimensions (economic, social and environmental) of sustainable development. This integrated perspective therefore demands, from the outset, research programmes and policy analysis that can establish how different aspects of the SDG agenda influence each

other, and whether actions taken to meet one specific SDG are likely to result in additional improvements towards other Goals ('co-benefits'), or to work antagonistically ('trade-offs').

As Nilsson et al. (2016) remark "Implicit in the SDG logic is that the goals depend on each other – but no one has specified exactly how." Indeed, the implied dependence between pairs of SDGs, or even between pairs of targets, depends on how precisely the underlying policy agenda is constructed and implemented. It therefore in general should be expected to vary from one country to another and to vary over time. The extent to which national policy enables co-benefits or inhibits trade-offs is therefore a measure of 'policy coherence' that is itself part of the SDG agenda: target 17.14 is precisely to "Enhance policy coherence for sustainable development". But the challenge of policy coherence is broader, since it shapes the entire worldview that the SDGs introduce; hence the calls for cohesive strategies, capacity-building initiatives and enabling environments for sustainable development that

E-mail address: J.H.P.Dawes@bath.ac.uk

appear, for example, in paragraphs 63 and 87 of the original Resolution adopted by the UN General Assembly (United Nations, 2015). Moreover, this is situated in the context of a yet wider debate on the nature of sustainable development and the extent to which it requires a global paradigm shift. This lies far beyond the scope of this paper but the dual policy ambitions to combine the achievement of minimum standards in social provision while ensuring that planetary boundaries are respected lies at the heart of the 'doughnut economics' movement (Raworth, 2017).

From the outset then, a study of interlinkages in the SDGs is therefore, at least to some extent, a study of policy coherence. Indeed, as discussed by Allen et al. (2018) the policy challenge set by sustainable development is to find policy interventions that are tangible, straightforward to implement, and which have high potential to effect transformational change. Many such interventions may have in fact a relatively low degree of transformational effect but nevertheless result in 'spillover' effects. Such spillovers can therefore be described in terms of either positive or negative influences on one or more of the SDGs in addition to the one for which the original policy action was designed. When these spillovers are positive the system demonstrates *policy coherence*: 'the systematic promotion of mutually reinforcing policy actions ... creating synergies towards achieving the defined objective' as defined by the OECD and quoted by Breuer et al. (2019).

Pham-Truffert et al. (2020) distinguish in their introduction between three levels of "SDG interaction": (i) interactions at the policy level resulting from conflicting global viewpoints and political priorities; (ii) interactions due to resource limitations; and (iii) systemic interactions that give rise to unintended consequences. They frame their study, and by implication the GSDR interlinkage matrix that they construct, and we study later in this paper, in terms of this third set of systemic interactions. But it seems difficult to draw precise boundaries between these different kinds of interaction: resource allocations are themselves policy decisions, and policy actions can be used to control unintended consequences, at least once the unintended consequence has come to light. It seems simpler instead to consider interlinkages as reflecting business-as-usual policy actions and to contrast this with the effects that novel, disruptive actions might lead to.

In practice, interlinkages under business-as-usual scenarios are more often inferred rather than explicitly set out in policy terms. To give just one example, Weitz et al. (2018) state that the construction of their cross-impact matrix of interactions was guided by the question "If progress is made on target x (rows), how does this influence progress on target y (columns)" (Weitz et al., 2018, page 536). The policy mechanism that would give rise to that linkage remains implicit. All methodologies for the identification of interlinkages proposed in the literature have drawbacks; in our examples of interlinkage matrices used later in this paper we focus on the use of expert analysis but this should not be taken as excluding interlinkage networks built from other data sources: the methodologies we propose would apply equally to these and we encourage other researchers to apply and build on these methods in their work. In particular Lusseau and Mancini (2019) point to two additional common and contrasting methodologies for the inference of interlinkages: historic analysis of time series of relevant data, for example World Bank indicators, or the identification of shared concepts in the underlying definitions of policy objectives.

There are at least four aspects to the identification of interlinkages; all four aspects themselves vary over time and geography, and all are challenging for the characterisation of a specific interlinkage: (i) its strength, (ii) its directionality (one way or in both directions), (iii) whether or not, over time, it is reversible, and (iv) the overall level of uncertainty (Nilsson et al., 2016). For example, the use of historical correlations in indicator timeseries might result in a reassuringly quantitative measure of the strength of an

interlinkage, while being less able accurately to determine its directionality or its variations in strength over time.

But such measures also need not be quantitative; the major contribution of the paper by Nilsson, Griggs and Visbeck is the proposal of their seven-point scale, from +3 ('Indivisible') to -3 ('Cancelling'), set out here in Table 1 for immediate reference, which provides a systematic scale, originally of a qualitative nature, around which evidence and decision-making can be organised. This scale was employed, for example, by the International Council for Science (known as ICSU after their original name: the International Council of Scientific Unions) in their 2017 report (International Council for Science (ICSU), 2017) that focussed on four SDGs (numbers 2, 3, 7 and 14).

In general one might hope to be able to draw on a sufficiently diverse collection of interlinkage networks, constructed independently, from which the most robust conclusions could be drawn, or differences between them be understood; this might, for example reveal how interlinkage strengths had evolved over time. Recent work in the area shows that this project has been initiated in the academic literature but it is far from concluded.

More fundamentally, the relation between the SDG interlinkage literature and policy studies for sustainable development also deserves more detailed scrutiny. In the public policy literature there is a central distinction between policy instruments (or tools, or techniques) and policy goals (outcomes). In terms of SDG interlinkages, the policy goals are made much more explicit than the policy instruments that might influence them; the mechanisms or inputs required to achieve the SDGs are often, as noted above, only described implicitly. The discussion of the relation between policy instruments and policy goals stretches back at least to the work of Tinbergen (1952) who proposed that if n independent policy goals were to be achieved, then a set of at least n policy instruments would in general be required (Tinbergen, 1952; Schaeffer and Willardsen, 2019). But of course, any degree of alignment or coherence between policy goals might enable a reduction in the number of policy instruments required. In the context of agroforestry and the SDGs, van Noordwijk et al. (2018) point to the meeting point of the SDGs in land use, where a small number of policy instruments could, at least in theory, help to achieve both environmental and resource extraction goals, as well as addressing issues of governance and inequalities. A set of seventeen separate policy instruments to achieve the SDGs would seem both unachievable and completely at odds with the call for policy coherence, yet the pursuit of some minimal number of policy 'tasks' does appear to lie at the root of various proposals, such as the Six Transformations proposed by the Sustainable Development Solutions Network (SDSN) (Sachs et al., 2019; Sachs et al., 2020) to which we will return in the final part of this paper, or the six 'Entry points' listed by the Global Sustainable Development Report 2019 Independent Group of Scientists appointed by the Secretary-General (2019). To sum up, when we refer to 'interlinkages between SDGs' we have in mind something closer to 'interactions between policy actions to address SDGs' although for brevity we

Table 1

The seven-point scale proposed by Nilsson et al. (2016) for scoring the influence of one specific SDG or target on another.

Score	Name	Explanation
+3	Indivisible	Inextricably linked to the achievement of another goal
+2	Reinforcing	Aids the achievement of another goal
+1	Enabling	Creates conditions that further another goal
0	Consistent	No significant positive or negative interactions
-1	Constraining	Limits options on another goal
-2	Counteracting	Clashes with another goal
-3	Cancelling	Makes it impossible to reach another goal

will continue to speak of 'SDG interactions' even though this is, in a very technical sense, a phrase that on its own lacks meaning.

We now introduce briefly the networks that we use as examples in the paper. Shortly after the launch of the SDGs, the International Council for Science and the International Social Science Council (ISSC) published an expert analysis of the formulation and structure of, and interlinkages between, the SDGs in 2015 (International Council for Science (ICSU), 2015). ICSU and ISSC merged in 2018 to form the International Science Council (ISC) but for historical consistency in this paper we will refer to the 2015 report (International Council for Science (ICSU), 2015) as the 'ICSU report'. The ICSU report (International Council for Science (ICSU), 2015; see also (Le Blanc, 2015)) analyses, from a 'science perspective' both the framing of the individual Goals, and targets within them, and the ecosystem of linkages between different Goals. The expert commentary suggests links between each Goal and specific targets within other Goals, and gives a sense of both the direction of influence and whether it is a reinforcing link or a negative trade-off between them. This report is particularly valuable for several reasons. Most importantly, while SDG 17 is not considered, the report does provide an expert analysis of SDGs 1–16 and the interlinkages between all of these, in a uniform manner. In the policy sense, then, these interlinkages are described in terms of the most typical and natural outcomes of policy implementation, and the report takes a global view in the knowledge that regional or national implementations may well depart from the policy agenda that is implicitly assumed in that report. The ICSU report is notable since few reports cover such a large subset of the SDGs; for example the later report (International Council for Science (ICSU), 2017) presents detailed analysis of interlinkages focusing on just four Goals: 2 (Zero Hunger), 3 (Healthy Lives), 7 (Energy) and 14 (Life Below Water), and the analyses of Blanc et al. (2017) and Singh et al. (2018) focus squarely on SDG 14.

A quantitative interaction network was derived from the expert analysis presented in the ICSU report previously by Dawes (2020). This quantitative version of the interlinkage network provides a clearer theoretical understanding of how the SDGs fit together, and allows system-level insights and conclusions to be drawn. At a system level, the work presented in Dawes (2020) showed that these SDG interlinkages promote faster progress on the first three SDGs (No Poverty, Zero Hunger, and Healthy Lives) compared to the remainder, and point to only very weak progress on Goal 15 (Life on Land) and indeed to negative progress on Goal 14 (Life Below Water), due to trade-offs in the system. The present paper extends this previous work through the comparisons with two other networks, the discussion of the sensitivity of these results to perturbations in the network, and the analysis of hierarchy between nodes that the network implies.

As an expert analysis, the ICSU report is complementary to the recent quantitative literature, much of which is based either on (historical) correlations in SDG indicators (e.g. Spaiser et al., 2016; Pradhan et al., 2017; Ranganathan et al., 2017; Ranganathan and Bali Swain, 2018; Bali and Ranganathan, 2021) or on Integrated Assessment Models, (e.g. Pedercini et al., 2019; van Soest et al., 2019). As we have noted above, all approaches have limitations; for example data-driven studies deriving from the analysis of historical correlations cannot uncover the directionality of influences, or underlying causation, e.g. the extent to which intended policy coherence directly resulted in the historically observed correlation. On the other hand, the ICSU report is limited by the relatively small group of authors, the need to find a consensus view that might end up not being applicable to any individual region or country, and the difficulty in quantifying precisely the strength of any interlinkage proposed.

The second expert analysis and literature review that we use as an example is the Global Sustainable Development Report (Independent Group of Scientists appointed by the Secretary-General, 2019) *The Future is Now: Science for Achieving Sustainable Development*, the first of the quadrennial reports produced for the UN's High-Level Political Forum on Sustainable Development by an independent group of scientists, while the third is a survey-based analysis of the coverage of Integrated Assessment Models (IAMs) carried out by van Soest et al. (2019) which we refer to as the 'IAM survey'. While in some aspects the most sophisticated of the kinds of model we consider, Integrated Assessment Models suffer from not being able to represent sufficiently well many of the socially-relevant SDGs, in particular SDG 5 (Gender Equality) and SDG 10 (Reduced Inequalities); we discuss this further in Section 4.

The GSDR 2019 Report contains a broader-based literature review carried out in early 2018, using keyword searches on the ISI Web of Science and Scopus databases, looking for articles that referred to both 'SDG' and 'interaction' or closely related terms (Pham-Truffert et al., 2020). From this analysis, the authors identified a collection of 65 global assessments (including UN reports) and a further 112 published scientific articles. The literature review extracted data at the target level (i.e. interactions between individual targets) which was then summarised at Goal level (see Box 1–2 on page 6 of Independent Group of Scientists appointed by the Secretary-General (2019)), and also in (Pham-Truffert et al., 2020, Fig. 4). This dataset therefore shows where the global scientific and science-policy communities have focused their attention to date across the Goals and particular interactions. The identification of the directionality and sign of interlinkages that are identified in this study do indicate how we expect a specific pair of SDGs to influence each other, but clearly the coverage of the network indicated by the GSDR 2019 report will reflect biases and non-uniformities in the scientific funding and research landscape, rather than properties intrinsic to the SDG network itself. This interlinkage data is summarised in Tables 2,3 in Appendix A for reference. A brief initial look shows, for example, that the influences of other SDGs on Goals 1–3 are more commonly reported than influences on later Goals, and that Goals 6 (Clean Water and Sanitation) and 7 (Affordable and Clean Energy) appear to be more influential than others.

A comparison between the interlinkages identified across these three reports (ICSU, GSDR 2019, and the IAM survey) allows us therefore to consider the robustness of the results reported in Dawes (2020) which were based on the ICSU report alone, as well as offering insight into possible gaps in the literature concerning particular interlinkages that we might expect to be significant. In this paper we explore both of these issues, as well as introducing relevant mathematical ideas to quantify three specific features of interlinkage networks that we see as useful in interpreting them for policy-related work. Firstly, the shape of the dominant 'mode of network response' that distinguishes which SDGs benefit most from the co-benefits in the network, and which are subject overall to the strongest trade-offs. Secondly, the sensitivity of key network properties to the addition of new interlinkages (or the variation in strength of existing linkages). Thirdly, the extent to which the directed links in the SDG network form a coherent overall hierarchy that allows the identification of some key 'enabling' or 'upstream' SDGs that drive progress on other 'downstream' Goals.

Our main conclusions are that there are strong similarities between the interaction matrices deduced from the ICSU and GSDR reports. Both have a single dominant 'mode' that shapes the overall system-wide response of the network. In the case of the GSDR report, the Goals most at risk of not being achieved include SDG 6 (Water), SDG 14 (Life below Water), SDG 15 (Life on Land), and SDG 17 (Partnerships for the Goals), see Figs. 2 and 4. There appear

to be structural similarities between the GSDR and ICSU Reports in that more progress is expected on SDGs 1, 2, and 3 than on others and there is a lack of attention as to how progress on Goals 1–3 feeds back into progress on Goals 4–17. The networks are therefore more sensitive to the addition, or strengthening, of interlinkages that reach from SDGs 1–3 to others, i.e. linkages in which progress on SDGs 1–3 is made in such a way that it explicitly opens up further progress on other SDGs. Similar messages emerge from the analysis of hierarchy within the SDG networks: SDGs 1 and 2 are consistency far ‘downstream’ of the other Goals, meaning that progress on SDGs 1 and 2 often reaps the rewards and benefits of progress elsewhere. In contrast, SDGs 12 (Sustainable Consumption and Production) and 17 (Partnerships) consistently appear far ‘upstream’, i.e. these are Goals that are enablers of progress. Similar conclusions are inferred from the IAM network, although, for reasons that we discuss below, these results contain a higher level of uncertainty than the ICSU and GSDR 2019 networks.

A final preliminary remark concerns the use of quantitative analysis in such a complex and multilayered policy-driven arena. The reduction of any set of SDG interactions to a matrix of numerical values will always inevitably be a hugely incomplete and highly reductionist viewpoint; it must be treated with caution. However, an important motivation for this kind of analysis is to develop a much greater sense of how a collection of individual interlinkages builds into a system-level representation that allows broader conclusions to be inferred. In this sense, then this paper attempts to provide tools to explore the consequences of one or other specific set of interlinkages, while acknowledging that the precise interpretation of any one set of results is likely to be affected by biases, often structural, that arise from the underlying data collection methodology. Attempting to make this connection between interlinkage inputs and system-level consequences is nevertheless important, even if the results are imperfect.

The structure of the paper is as follows. In Section 2 we set out the construction of interlinkage networks from the ICSU and GSDR 2019 reports, and in Section 2.2 we provide a detailed discussion of eigenvalues and eigenvectors and their relationship to concepts of centrality in networks. This motivates their applicability here and aids the interpretation of the results in Section 2.3 for the ICSU and GSDR interaction matrices. In Section 2.4 we use the separate totals of positive and negative influences summarised in the GSDR 2019 report to test the robustness of the leading eigenvalue and the shape of the leading eigenvector to random sampling over these positive and negative values. Section 3 defines two distinct sensitivity measures for the interlinkage matrices, depending on whether one wishes to improve the overall growth rate of the dominant mode of response (i.e. to intensify the influence of the interactions), or to equalise the progress described by the dominant mode across all the SDGs. In Section 4 we briefly summarise similar results on the leading eigenvalue and eigenvector for the IAM network, derived from the results of an expert survey on interlinkages related to Integrated Assessment models. These results are perhaps less reliable than those presented in Section 2 but are nevertheless consistent with them. Section 5 considers the question of overall hierarchies within these weighted, directed networks, for all three interaction matrices. Finally, Section 6 presents conclusions and directions for future work.

2. SDG interlinkages

2.1. Network structures

In this section we review the construction and interpretation of interlinkage networks derived from the ICSU 2015 report (International Council for Science (ICSU), 2015) and the GSDR

2019 report (Independent Group of Scientists appointed by the Secretary-General, 2019). We then compare the results and present a preliminary discussion on the implications for where this might indicate gaps in the published literature. It is important to note that the ICSU 2015 report omitted Goal 17 (Partnerships for the Goals) and discussed interactions between Goals 1–16 only.

2.1.1. ICSU network

The methodology used to construct a network from the expert analyses presented in the ICSU 2015 report is presented in detail elsewhere (Dawes, 2020, section 2.1). In brief, the expert analysis proposed linkages between Goals, in effect estimating the effects that typical policies to achieve one goal would have on targets within other SDGs, implicitly using the viewpoint discussed in Section 1. This expert analysis allowed such linkages to be proposed between each Goal and specific targets within each of the other Goals. Narrative descriptions of these linkages usually allowed a direction of influence to be inferred (in some cases mutual co-benefits, or trade-offs, which were represented by bidirectional linkages), with an interaction strength given by the number of specific targets indicated as a proportion of the total number available (note that the number of targets per Goal varies between Goals). In this fashion a set of weighted, directed interlinkages at the level of entire Goals, can be deduced directly from the narrative in the report. In the resulting adjacency matrix A , the influence of SDG j on SDG i is described by the matrix entry A_{ij} . Each linkage can potentially be identified and described twice, in the narrative commenting on each of the Goals it connects. This is accounted for, and the interaction strengths are scaled to lie in the range $-1 \leq A_{ij} \leq 1$.

The ICSU adjacency matrix A is illustrated in Fig. 1(a). Since self-reinforcing links (i.e. from a Goal to itself) were not allowed, the matrix has zero entries on the main diagonal. Of the remaining 240 entries, 162 are non-zero: 152 of these are positive and 10 are negative. Since many of these are close to zero, the illustration in Fig. 1(a) colours white the entries that are close to zero, as indicated by the colour bar to the right of the figure. In terms of significant entries, there are 46 cases in which $A_{ij} \geq 1/3$ and 2 cases in which $A_{ij} \leq -1/3$. The two significant negative entries are $A_{14,2}$ and $A_{14,11}$: the influence of SDG 2 (Zero Hunger) and SDG 11 (Sustainable Cities) on SDG 14 (Life Below Water).

It is particularly interesting to note that there is a large mainly white area in the leftmost three columns of the matrix, showing that there are few influences of SDGs 1–3 on the later Goals 4–16. In contrast, the top three rows of the matrix contain a large number of significant positive entries shown in brighter (yellow and orange) colours. These indicate that progress on SDGs 4–16 in many cases drives progress on SDGs 1–3. Within the collection of SDGs 4–16 there are a large number of interactions, most of these positive showing the overall self-reinforcing nature of the SDGs viewed at a system level.

2.1.2. GSDR network

Turning to the GSDR 2019 report (Independent Group of Scientists appointed by the Secretary-General, 2019), a similar interaction network was computed by these authors based on their literature survey of 177 global scientific assessments, UN flagship reports and scientific articles on interlinkages between the SDGs. These 177 reports and journal articles were analysed, wherever possible, at the level of individual targets, as described by Pham-Truffert et al. (2019) and Pham-Truffert et al. (2020). Articles were read by hand in order to identify statements in which authors indicated specific links between targets (or complete SDGs) and assessed on the seven-point scale proposed by Nilsson et al. (2016). Directionality was also inferred from the text. Results were aggregated, keeping positive and negative scores separate; as the




Fig. 1. (a) ICSU interaction matrix, reproduced from (Dawes, 2020, Fig. 1). (b) The averaged interaction matrix produced by the GSDR 2019 report (Independent Group of Scientists appointed by the Secretary-General, 2019), showing the average interaction strength $A_{ij}^{av} := (N_{ij}^+ - N_{ij}^-)/N_{max}$ for each link $j \rightarrow i$. Blanks indicate values close to zero.

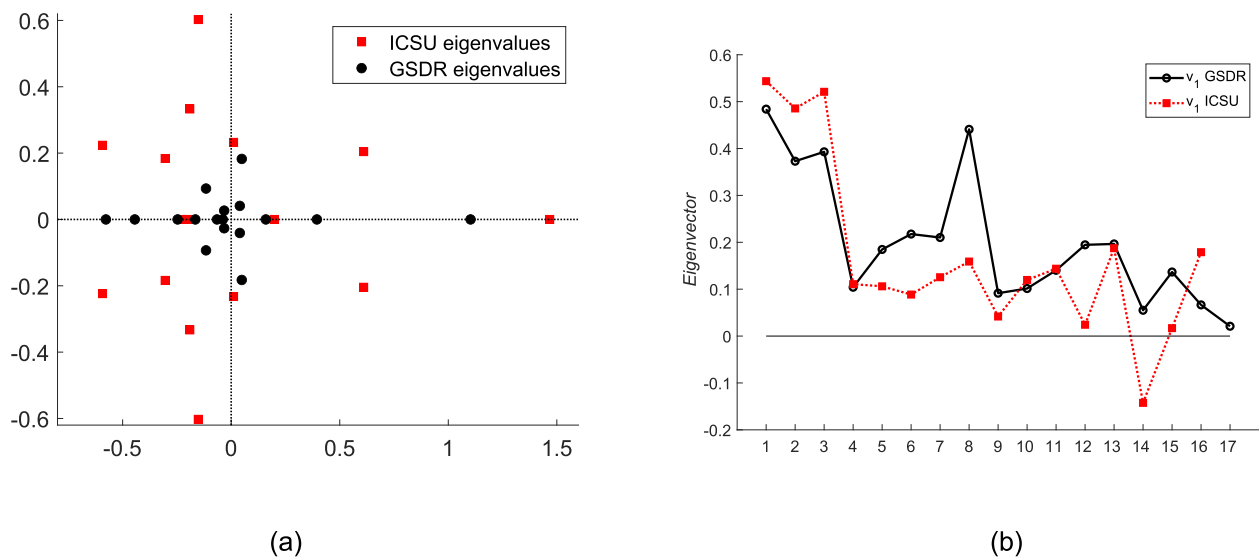


Fig. 2. Comparisons between the interlinkage matrices from the ICSU Report (International Council for Science (ICSU), 2015) and the GSDR 2019 Report (Independent Group of Scientists appointed by the Secretary-General, 2019). (a) Eigenvalues of the (averaged) interlinkage matrices. (b) Components $v_1^{(1)}, \dots, v_{16}^{(1)}$ of the eigenvector $\mathbf{v}^{(1)}$ corresponding to the eigenvalue λ_1 having the largest real part, in the two cases. The horizontal axis indicates the SDGs by number. There is no component for SDG 17 shown on the red dotted line as the ICSU network considered only SDGs 1–16. Each eigenvector is normalised so that the root-mean-square of the entries is equal to one.

GSDR authors note (Pham-Truffert et al., 2020) the practice of combining positive and negative scores to produce only a net score often obscures the potential for negative influences. The analysis of the development literature carried out for the GSDR 2019 report appears implicitly to deal with policy as it appears to be currently implemented rather than any more ambitious or forward-looking policy design that might transform outcomes.

The use of the seven-point scale is also worth commenting on. For a single research report, or policy initiative it feels entirely reasonable to try to score an interaction on the seven-point scale (see Table 1), in order to aid decision-making and force analysts to come to some kind of agreement over, for example, the relative importance of different influences between SDGs. When attempting to aggregate the effect of many possible influences, it is not clear that one should simply sum up, or average over, the influence

scores from several sources. For example, two separate policy choices may both result in a positive influence of one SDG on another yet be themselves mutually exclusive due to resource constraints. In this case the set of options should ideally not include both influence scores since an ‘either-or’ choice is required in policy terms. Despite shortcomings of this kind, the use of the scale has many advantages, not least in the imposition of clarity both in terms of evaluation of the set of possible policy actions, and the methodology by which these quantitative scores are computed.

Although the GSDR data is presented at target level, in this paper we will consider only the aggregated data that they present at the level of entire Goals since, as the GSDR authors note, the target-level coverage is considerably more sparse. The GSDR literature review produced a total of 5,758 interactions at the Goal level: 4,976 positive influences and 782 negative. The numbers of

positive and negative interactions are presented in Tables 2 and 3 respectively. These tables also contain the row and column sums for each Goal which capture the total number of influences of other Goals on a particular Goal (the row sum, known as the 'in-degree' of the network node) and influences of a particular Goal on all the others (the column sum, known as the 'out-degree' of the network node).

In considering interactions only at the level of entire Goals, aggregating across targets within each Goal, the diagonal entries in the adjacency matrix from the GSDR 2019 data lose some of their meaning: instead of demonstrating that targets within the same Goal are linked, they represent just that progress overall on a Goal would lead to further progress. For consistency, both within the GSDR network, and in comparison to the ICSU and IAM networks, neither of which (by construction) has diagonal entries, we set the diagonal entries in the GSDR matrix to zero. The effect of this change is small, as is shown by Fig. 11 in the Appendix where we plot the eigenvalues and leading eigenvector for the GSDR matrices with and without the diagonal elements.

For each pair of Goals (i, j) we define N_{ij} as the number of sources that describe a directed linkage from Goal j to Goal i . Due to multiple linkages being possible at the target level, and the subsequent aggregation of target level data up to Goal level, the largest number of sources for a single linkage from one SDG to another is slightly larger than 177 (the number of sources used the literature review): $N_{\max} = 184$. This occurs for the linkages from Goal 7 to Goal 6 where the GSDR 2019 Report identifies 123 positive linkages and 61 negative linkages between (targets associated with) these two Goals. Further, we denote the number of positive, reinforcing links where progress on Goal j reinforces progress on Goal i , by N_{ij}^+ . Similarly, the number of links with a negative weight, indicating a trade-off between Goal j and Goal i , is denoted by N_{ij}^- . The total number $N_{ij} = N_{ij}^+ + N_{ij}^-$.

Mindful of the potential issues involved in aggregating positive and negative scores, as discussion in the second paragraph of this section above, we define the adjacency matrix A^{av} for the GSDR data to be the average interaction strength, i.e. the difference between the total positive and negative contributions, scaled by N_{\max} . Precisely, we define $A_{ij}^{\text{av}} := (N_{ij}^+ - N_{ij}^-)/N_{\max}$; the numerator corresponds to the entries shown in Table 4 for reference. The adjacency matrix A^{av} is illustrated in Fig. 1(b). By construction we again have $-1 \leq A_{ij}^{\text{av}} \leq 1$, for each pair (i, j) , as for the ICSU case.

The overwhelming majority of matrix entries, 228 out of a possible 272 (i.e. for the 17×17 matrix but with the diagonal entries removed), are positive, described in the GSDR report as producing 'co-benefits to be harnessed'. In addition, 28 matrix elements are zero (excluding the diagonal entries), and 16 are negative, described as 'trade-offs to be addressed'. In Fig. 1(b) where cells are again coloured white if their values do not significantly differ from zero, we see that none of the negative entries is less than -0.1 , and indeed only four are less than -0.025 . These four are: $A_{6,1}^{\text{av}} = -0.10$, $A_{15,1}^{\text{av}} = -0.054$, $A_{14,13}^{\text{av}} = -0.098$, and $A_{7,15}^{\text{av}} = -0.054$.

Although the interlinkages described by the ICSU Report and the GSDR 2019 Report differ in numerous ways when considered in detail; it is interesting to note similarities in the form of the adjacency matrix. The GSDR matrix in Fig. 1(b) has greater white space in the lower left of the figure compared to the top three rows, similar to the ICSU matrix. SDGs 4, 5, 9, 10, 16 and 17 have noticeably lower in-degrees than the other SDGs; this is highlighted by the amount of white space in each of these rows in Fig. 1(b). Similarly, SDGs 1, 3, 4, 5, 10 and 16 have noticeably lower out-degrees than the others; there is more white space (or light blue shading) in these columns in Fig. 1(b). These results relate directly to the discussion in Pham-Truffert et al. (2020) classifying SDGs as 'buf-

fers' (where the in-degree of the node is higher than the out-degree), or 'multipliers' (where the out-degree of the node is higher than the in-degree). Interestingly, the ICSU interaction matrix also shares these in-degree and out-degree characteristics for these SDGs, with the exception that the out-degrees for SDGs 4, 5 and 10 do not appear to be significantly lower than for other Goals.

2.2. Centrality measures, eigenvalues and eigenvectors

In this section we review the conditions for the two adjacency matrices to show 'self-consistent' behaviour, and investigate whether these conditions hold. Given the significant number of negative linkages in the GSDR dataset, we also investigate the distribution of possible adjacency matrices that these could represent, and the properties of this wider distribution of networks; this provides a sense of the robustness of these results for the GSDR network.

To provide wider context, we begin by recalling the general idea in network science of the 'centrality' of a node in a network: centrality measures attempt to make a quantitative estimate of the relative importance of different nodes (Newman, 2018). The most obvious notion of centrality is the number of neighbours d_i that a given node i has - this is known in the literature as *degree centrality*. For a directed network the numbers of ingoing and outgoing edges connected to node i can be calculated separately; the results are the in-degree k_i^{in} and out-degree k_i^{out} for each node i . In terms of our adjacency matrix A where the element A_{ij} represents the strength of an interlinkage from node j to node i , we can write these centrality measures by summing over one of the indices of the matrix A :

$$k_i^{\text{in}} = \sum_{j=1}^n A_{ij}, \quad k_i^{\text{out}} = \sum_{j=1}^n A_{ji}, \quad k_i = k_i^{\text{in}} + k_i^{\text{out}}. \quad (1)$$

To be precise, since the entries in A are weighted, rather than just being either 1 or 0 to show the presence or absence of an edge, these measures define the *weighted in-degree*, *weighted out-degree*, and *total weighted degree*, respectively. These are the measures that underpin the analysis by Pham-Truffert et al. (2020), for example, and from which they describe their typology of roles played by different nodes, including the notion of 'buffers' where k^{in} is much larger than k^{out} and 'multipliers' where the reverse is true.

One criticism of the use of (weighted) in-degree and out-degree as centrality measures is that these measures take account only of the local connectivity of each node. Other centrality measures, such as eigenvector centrality, or 'eigencentrality', take a more global view of the network. Eigencentrality is motivated by defining the importance v_i of node i in terms of the importance of the nodes that it is connected to, weighted by the strength of those connections, i.e.

$$v_i := \frac{1}{\lambda} \sum_{j=1}^n A_{ij} v_j, \quad (2)$$

where λ is an overall weighting parameter. At first sight this equation appears to be distinctly unhelpful since it is entirely self-referential: in order to compute v_i for node i I need to know v_j for all nodes j that connect to i . However, multiplying up by λ and considering all n nodes together we see that (2) is equivalent to the matrix-vector equation $\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$ which is mathematically well understood. Generically an $n \times n$ matrix A will have n distinct eigenvalues $\lambda_1, \dots, \lambda_n$ and a corresponding collection of eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ each satisfying $\mathbf{A}\mathbf{v}_i = \lambda_i \mathbf{v}_i$ for $i = 1, \dots, n$. If the entries of A are all non-negative then the largest (also known as the 'leading') eigenvalue λ_1 of A is guaranteed to be real and to have an eigenvector (the 'leading eigenvector') \mathbf{v}_1 that has all components of the

same sign (or zero); this is the Perron–Frobenius Theorem (Newman, 2018). The elements of the vector \mathbf{v}_1 then satisfy (2) with $\lambda = \lambda_1$ and hence can be interpreted as a measure of the relative importance of each node in the network; they are known as the *eigencentralities* of the nodes. We note that the eigencentralities of nodes are defined only relative to the eigencentralities of other nodes since the eigenvectors themselves are defined only up to an overall scale factor.

Eigencentrality has a long history within the social networks literature, stretching back at least as far as the work of Bonacich (1972). As Bonacich (2007) and Bonacich (2007) notes, eigencentrality takes account not only of the direct connections to and from each node, but also the connectivities of these neighbours, and the neighbours of neighbours, and so on. It is therefore a centrality measure that takes account of the influence of a node across the entire network structure.

The eigencentrality measure can be related to the definition of *self-consistency* for the network with adjacency matrix A proposed by Dawes (2020): in that paper a network was defined to be *self-consistent* if λ_1 is real and positive and has an eigenvector \mathbf{v}_1 which has all entries positive. If all entries of A are themselves positive then A is self-consistent by the Perron–Frobenius Theorem, but this is not a necessary condition for A to be self-consistent, i.e. the set of self-consistent adjacency matrices is larger than the set of positive adjacency matrices.

As well as defining eigencentrality, there is another sense in which the leading eigenvalue and associated eigenvector of A is important, which we turn to now; this is motivated by a dynamical view of the interlinkages as encapsulated by the statement in Weitz et al. (2018, page 536), noted earlier: “If progress is made on target x (rows), how does this influence progress on target y (columns)”.

The simplest dynamical model for the intrinsic interactions between the SDGs would be the linear differential equation $d\mathbf{x}/dt = A\mathbf{x}$ where $\mathbf{x}(t) = (x_1(t), \dots, x_n(t))$ is the vector of values describing the level of progress $0 \leq x_i(t) \leq 1$ on SDG i at time t , rescaled to lie in the interval $[0, 1]$, i.e. to represent the proportion of the target achieved so far, as discussed in Dawes (2020). This is of course an extremely simple dynamical model, and far more elaborate modelling has been carried out by several groups (Spaiser et al., 2016; Ranganathan et al., 2017; Ranganathan and Bali Swain, 2018; Pedercini et al., 2019). The simple linear equation can be justified on grounds on minimality and that over relatively short times, generically every dynamical system can be approximated by a linear one.

Of particular relevance here is that the solutions to $d\mathbf{x}/dt = A\mathbf{x}$ are in many cases described extremely well by just the leading eigenvalue λ_1 and its corresponding eigenvector \mathbf{v}_1 . This statement holds in particular in the case of a self-consistent matrix A for which the leading eigenvalue λ_1 is real and (much) larger than the real parts of all the other eigenvalues $\lambda_2, \dots, \lambda_n$, i.e. so that $Re(\lambda_n) \leq Re(\lambda_{n-1}) \leq \dots \leq Re(\lambda_1)$, where n is the size of the adjacency matrix and Re denotes the real part of the (possibly complex) eigenvalue. The interpretation is that for such a matrix the exponential growth described by $d\mathbf{x}/dt = A\mathbf{x}$ is dominated by the leading eigenvector \mathbf{v}_1 as it is the dominant ‘mode of response’ of the system. The amplitude of this ‘mode of response’ grows $\sim \mathbf{v}_1 e^{\lambda_1 t}$ which is faster than the remaining $n - 1$ modes.

The eigenvector corresponding to the leading (i.e. largest) eigenvalue therefore has a dual interpretation: for a matrix with non-negative entries it describes the eigencentrality of nodes, and for matrices in general it describes the components of the dominant mode of response of the linear dynamical system $d\mathbf{x}/dt = A\mathbf{x}$. In this way the leading eigenvector captures both a static sense of the relative importances of different nodes and

the dynamic sense of how the network nodes may reinforce, or counterbalance, each other over time.

2.3. Leading eigenvectors for interlinkage matrices

The interlinkage matrix A^{av} derived from the GSDR network is an example of a self-consistent matrix. Self-consistency (a leading eigenvalue that is real and positive) implies that the behaviour of the model $d\mathbf{x}/dt = A\mathbf{x}$ will show monotonic (i.e. not oscillating) exponential growth in solutions as time t increases. This is clearly a natural and desirable quality in interlinkage networks describing SDG influences.

Fig. 2 summarises these similarities and differences by showing, in Fig. 2(a) the set of eigenvalues of the two matrices. In both cases the eigenvalue with the largest real part is actually real (rather than being a complex conjugate pair of eigenvalues) and positive, indicating that the co-benefits and trade-offs overall lead to monotonic (and exponential) growth of progress on the SDGs. For the ICSU network, $\lambda_1 = 1.4672$; for the GSDR network $\lambda_1 = 1.1026$ (both to 4dp). Although the leading eigenvalues for the two matrices differ substantially, this is more a consequence of their different constructions and the distribution of entries and is not in itself important. Of more importance, and another point of similarity between the two cases, is that there is a noticeable gap between the leading eigenvalue λ_1 and the eigenvalue λ_2 with the next-largest real part. For the ICSU network we find $Re(\lambda_1 - \lambda_2) = 0.86$, and for the GSDR network $Re(\lambda_1 - \lambda_2) = 0.71$ (to 2dp). This indicates that the behaviour will in both cases be dominated by just one ‘mode’, given by the eigenvector corresponding to the leading eigenvalue.

Fig. 2(b) compares, for the ICSU and GSDR cases, the components of the (normalised) eigenvectors $\mathbf{v}^{(1)}$ corresponding to the leading eigenvalues λ_1 . In both cases the first three components, corresponding to SDGs 1–3 on No Poverty, Zero Hunger, and Healthy Lives, are significantly higher than almost all the remaining components. Both eigenvectors have low components for Goal 14 (Life Below Water). For the ICSU Report, this component is negative, meaning that over time progress on Goal 14 is likely to worsen rather than improve. For the GSDR 2019 Report, this component is positive (so the network implied by the GSDR 2019 Report is ‘self-consistent’ in the terminology of Dawes (2020) and summarised above) but it is small; indeed it is the smallest component apart from that corresponding to Goal 17 on Partnerships (note that SDG 17 was not included in the ICSU analysis). The high values for the components related to SDGs 1, 2, 3 and 8 are related to their roles as ‘buffers’ in the analysis of Pham-Truffert et al. (2020): there are positive ‘co-benefits’ reinforcing progress on each of these SDGs from several others.

As well as these two important similarities, there is one central difference: the much larger component for Goal 8 (Growth) which relates to the fact that SDG 8 has the fourth highest in-degree, after SDGs 1, 2, and 3 (which have the highest in-degrees, in fact in the same order as the numbering of the Goals themselves). This corresponds to row eight in the interaction matrix in Fig. 1(b) which contains significant positive entries in positions (8, 3), (8, 6), (8, 7) and (8, 15) showing that progress on Goal 8 is supported in particular by progress on the Goals related to Healthy Living (SDG 3), Water (SDG 6), Energy (SDG 7), and Life on Land (SDG 15), respectively.

2.4. Testing robustness by constructing an ensemble of GSDR networks

The appearance of only 16 negative entries in A^{av} (and only four that are significantly negative) obscures the fact that for 100 out of

the 272 possible matrix entries some fraction of the linkages identified are in fact negative. Negative entries are important since they point to trade-offs between SDGs which often demand additional attention for Agenda 2030 to be able to succeed overall. Moreover, the results presented by the GSDR depend crucially on the population of reports that are used to compile the data, and if a different set of reports had been analysed, the number of negative linkages might have varied due to the inherent variability in the possible sub-samples of the literature. Further, it seems likely that some of the negative effects become more prominent in particular settings; 'worst-case scenarios' may include cases in which the co-benefits are not available simply to cancel out trade-offs in the manner suggested just by summing all contributions in the construction of the GSDR network A^{av} above. To investigate the level of variation, we extend the analysis of the GSDR interaction matrix A^{av} by generating a distribution of interaction matrices in order to explore the robustness of the results presented above.

We generate an ensemble of 10^5 matrices A by choosing positive or negative values, independently for each entry A_{ij} , for all entries where $N_{ij} > 0$. If $N_{ij} = 0$ then we set $A_{ij} = 0$; but if $N_{ij} > 0$ then we set A_{ij} to take the.

- positive value $A_{ij} = N_{ij}^+ / N_{max}$ with probability $p_{ij} := N_{ij}^+ / N_{ij}$, or
- negative value $A_{ij} = -N_{ij}^- / N_{max}$ with probability $1 - p_{ij} = N_{ij}^- / N_{ij}$.

We then compute the eigenvalues of A and the eigenvector $\mathbf{v}^{(1)}$ for the leading eigenvalue λ_1 . The results from this ensemble of matrices are shown in Figs. 3 and 4. Fig. 3(a) illustrates the distribution of eigenvalues of these sampled matrices in the complex plane. Each matrix in the ensemble typically has a real leading eigenvalue λ_1 and a complex conjugate pair of eigenvalues λ_2, λ_3 that have the next-most-positive real parts. We find that λ_1 typically lies in the range $0.7 \leq \lambda_1 \leq 1.5$; for comparison recall that for A^{av} , $\lambda_1 = 1.1026$ which is close to the mode of the distribution for λ_1 shown in Fig. 3(b). The real parts of the complex conjugate pair λ_2, λ_3 vary approximately in the range 0.6–0.8, preserving the separation between them and λ_1 . Their imaginary parts vary roughly between zero and 0.4. Fig. 3(a) shows that the remaining eigenvalues $\lambda_4, \dots, \lambda_{17}$ are distributed further to the left, roughly within a circle of radius 0.5 centered on the origin as indicated by the blue-green areas in the plot.

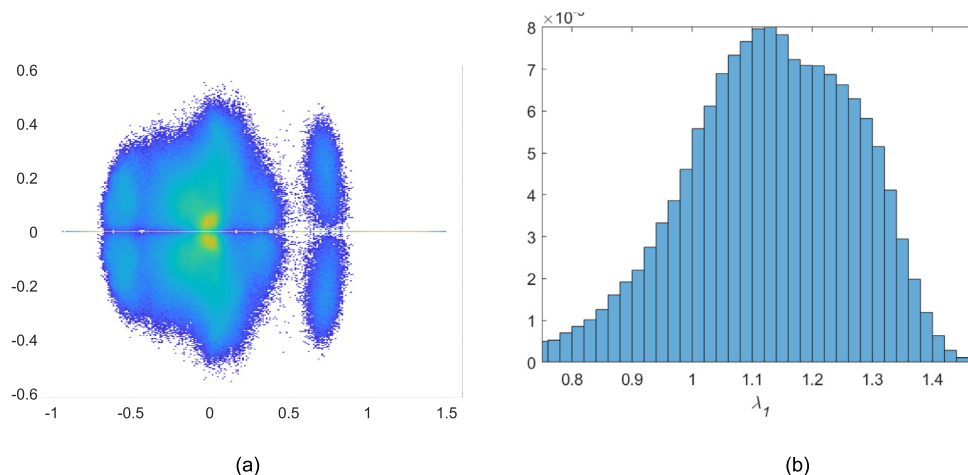


Fig. 3. Density plots of the locations in the complex plane of eigenvalues $\lambda_1, \dots, \lambda_n$ of the ensemble of GSDR interaction matrices from 10^5 samples of positive and negative interaction weights. (a) Heatmap showing density of all eigenvalues in the complex plane (yellow - high, dark blue - low). Note the yellow regions, indicating high density of an eigenvalue, on the positive real axis as well as near zero. (b) Histogram showing the distribution of the leading eigenvalue λ_1 along the real axis. This corresponds to the yellow-blue line along the real axis in (a).




Fig. 4. Violin plot showing the variation in components $v_1^{(1)}, \dots, v_{17}^{(1)}$ of the eigenvector $\mathbf{v}^{(1)}$ corresponding to the leading eigenvalue λ_1 of the GSDR interaction matrix. The horizontal axis indicates the SDGs by their numbers. Note the significant shaded areas below the horizontal line at zero for Goal 6 (Water), Goal 14 (Life Below Water) and Goal 15 (Life on Land). For clarity of presentation, extreme outliers from the lowest and highest 0.5% of the simulation ensemble have been removed from the data shown in this figure.

Fig. 4 summarises the distribution of components of the leading eigenvector corresponding to λ_1 from matrices in the ensemble. These results are presented as a 'violin' plot (Hoffmann, 2015) in which the width of the shape indicates the relative frequency of the value, and the mean value for each component is indicated by the horizontal black bar. In many cases the distributions are multi-humped, with two or three separate maxima, leading to repeated expansions and contractions of the shapes in the plot (which are sometimes 'violin-shaped', hence the name). These fluctuations are caused by the presence of a relatively small number of high-weight edges in the network that contain significant negative contributions; as these switch sign, the shape of the eigenvector changes significantly, meaning that the distribution is quite strongly influenced by a relatively small number of interlinkages rather than being a composition of many independent interlinkages of more equal sizes.

In Fig. 4 it is particularly interesting to note that some SDGs are associated with much greater variability than others. In particular, SDGs 1, 3, 6, 14 and 15 have elongated distributions while the distributions for SDGs 4, 5, 9, 16 and 17 are much more compact. Fur-

ther, SDGs 6 (Water), 14 (Life Below Water), and 15 (Life on Land) have significant part of their distributions lying below zero. These parts of the distributions, which signify negative progress on these Goals, indicate that these three SDGs are most at risk. For comparison, and ease of interpretation, conventional histograms of the individual components of the leading eigenvector for the matrices in the ensemble are shown in Fig. 12 and Fig. 13 in Appendix A.

3. Sensitivities

Since the nature of policy actions, and therefore implicitly these typical SDG interlinkages, vary across regions of the world, and indeed from country to country, a natural set of questions arise as to where, if one were able through policy actions, the interlinkage network could be most usefully adjusted in order better to achieve the whole of Agenda 2030 overall. First, of course, it is necessary to define clearly what one would mean by the most valuable adjustments, or perturbations, to make to the SDG interlinkages. This, therefore, is a question of the sensitivity of the initial network to perturbations in each of the directed edges. What improvements at the system level can we make by adjusting individual links?

Of the many possible definitions of an ‘improvement’ to the SDG network, we focus in this section on two: we seek either to increase the multiplier effect associated with the leading eigenvalue of the adjacency matrix, i.e. to increase the growth rate through which the SDGs, allow self-reinforcing progress; or, we seek to make this progress more equal across the SDGs, i.e. to equalise the components of the leading eigenvector displayed in Fig. 2(b).

The mathematical and engineering literature on the effects of perturbations on the eigenvalues and eigenvectors of a matrix is large, extending back to the 19th century. The mathematical history, together with explicit results that we use here, has been reviewed recently by Greenbaum et al. (2020); in the engineering literature a foundational reference is the paper by Nelson (1976).

A suitable mathematical measure of the sensitivity of the growth rate to perturbations is to consider the rate of change of the leading eigenvalue λ_1 of the interlinkage matrix to perturbations in one element A_{ij} of the adjacency matrix. We therefore define the ‘multiplier effect’ sensitivity matrix S^m to be the $n \times n$ matrix whose (i, j) entry is $S_{ij}^m := (1/\lambda_1) \partial \lambda_1 / \partial A_{ij}$. The factor of $1/\lambda_1$ is included so that S^m computes relative changes in λ_1 rather than absolute changes. In the cases of interest here, S_{ij}^m turns out to have a simple expression in terms of the left and right eigenvectors ($\mathbf{y}^{(1)T}$ and $\mathbf{v}^{(1)}$, respectively) for the leading eigenvalue λ_1 . A straightforward calculation, presented in Appendix B for completeness, leads to

$$S_{ij}^m = \frac{1}{\lambda_1} \frac{y_i^{(1)} v_j^{(1)}}{\mathbf{y}^{(1)T} \mathbf{v}^{(1)}} \quad (3)$$

i.e. the matrix S^m is the outer product of the vectors $\mathbf{y}^{(1)}$ and $\mathbf{v}^{(1)}$, normalised.

Fig. 5 plots the growth rate sensitivity measure S_{ij}^m for each edge $j \rightarrow i$ in the interlinkage network, for both the ICSU and GSDR networks. It is immediately clear that perturbing different links in the network can lead to significantly different outcomes. The form of S^m given in (3) shows that, for a fixed value of the row index i , the values in each row have the same shape as the entries in the right eigenvector $\mathbf{v}^{(1)}$ which is shown in Fig. 2(b). Hence, reading across each row of Fig. 5(a) follows the rises and falls in the components of the leading eigenvector for the ICSU matrix, and similarly with Fig. 5(b) and the GSDR eigenvector from Fig. 2(b). For rows i where the left eigenvector $\mathbf{y}^{(1)}$ (not shown) has a negative component, the corresponding row in Fig. 5 has large negative

components where $\mathbf{v}^{(1)}$ is large and positive, for example in row 2 in Fig. 5(a).

The perturbations that increase the leading eigenvalue λ_1 most rapidly are, for the ICSU matrix, seen to be those that link from SDGs 1, 2 and 3 to Goals 4, 5, 9, 12 and 14. Each of these links is one that connects an SDG where significant progress would be predicted, to an SDG where progress is less well supported by the internal network structure. For the GSDR matrix, the links that most increase λ_1 are from SDGs 1, 2, 3 and 8 to Goals 6, 7 and 12. The fact that links from SDG 8 are favoured is connected to the large component on Goal 8 in the leading eigenvector, which is a notable difference in the ICSU and GSDR results, as discussed earlier. In the GSDR analysis, there are strong reinforcing links identified between SDGs 6 and 12 which explain why rows 6 and 12 in Fig. 5(b) are extremely similar. Links from SDGs 1, 2, 3 or 8 to SDGs 6 or 12 could be achieved through strategies that explicitly link progress on the first three SDGs to sustainable water resource use, recycling, and circular economy ambitions. Conversely, it is interesting to observe that the analysis does not support strategies to connect progress on SDGs 1, 2 and 3 to progress on SDGs 4 (Education), 16 (Peace) or 17 (Partnerships for the Goals). From the analysis presented in Fig. 5 it appears that enhancing those linkages within either SDG network would lead to margin results at best.

The alternative measure of the effectiveness of a perturbation to A is slightly more complicated to define, mainly because eigenvectors themselves are only defined up to a scalar multiple. This means we have to choose an appropriate normalisation of the eigenvectors in order to describe the rate of change of an eigenvector with respect to a perturbation to the matrix. Optimal equal progress would be made on all SDGs if the leading eigenvector had all components equal, i.e. $\hat{\mathbf{n}} := (1, \dots, 1)/\sqrt{n}$ where the factor of $1/\sqrt{n}$ is included as a normalisation, i.e. so that $\sum_{i=1}^n \hat{n}_i^2 = 1$. For the right and left eigenvectors $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(n)}$ and $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n)}$ respectively, we normalise by requiring $\|\mathbf{v}^{(j)}\| = 1$ and $\mathbf{y}^{(i)*} \mathbf{v}^{(i)} = 1$ for all i and j , where the asterisk denotes the complex conjugate transpose.

We then define the sensitivity S_{ij}^{eq} of the leading eigenvector $\mathbf{v}^{(1)}$ to be

$$S_{ij}^{\text{eq}} := \hat{\mathbf{n}} \cdot \frac{\partial \mathbf{v}^{(1)}}{\partial A_{ij}}. \quad (4)$$

That is, the entries of the matrix S^{eq} measure the extent to which a perturbation to A_{ij} increases the alignment of the leading eigenvector $\mathbf{v}^{(1)}$ with the vector $\hat{\mathbf{n}}$ that describes complete equality of progress on all SDGs. Mathematical details concerning (4) are given in Appendix B.

Interestingly, the two parts of Fig. 6 show overall similarities to the parts of Fig. 5, respectively. This can be explained in part by similarities in the mathematics, and the appearance of the product $\bar{y}_i^{(k)} \bar{v}_j^{(k)}$ in (11), similar to the term $y_i^{(1)} v_j^{(1)}$ in (3). But there are clear differences, for example for the ICSU network Fig. 6(a) shows that a strengthening of the interlinkages from SDGs {1, 2, 3} to SDG 14 would serve to equalise the leading eigenvector, but as indicated in Fig. 5(a) it would increase the leading eigenvalue only marginally.

Fig. 6 indicates that, in general, increasing the directed interlinkages from SDGs 1, 2 and 3 to other Goals (especially SDGs 5, 6, 9, 12, 14, 16 or 17) leads to a more balanced leading eigenvector, and so is more likely to lead to consistent progress across all the Goals, rather than the unbalanced progress indicated in Fig. 2(b) and Fig. 4 which emphasises progress on SDGs 1, 2 and 3. The intuitive explanation for these results is entirely in line with Fig. 2(b); if the higher progress on SDGs 1, 2, and 3, implicit within the interlinkage matrix, can be directed towards advancing progress elsewhere, particularly on SDGs that are less well supported by the




Fig. 5. Sensitivity matrices S^m for the leading eigenvalues λ_1 of (a) the ICSU network illustrated in Fig. 1(a); and (b) the averaged GSDR interaction network, illustrated in Fig. 1(b). Yellow and orange entries in the matrix indicate where increases in those interlinkages would increase the leading eigenvalue λ_1 ; dark blue entries indicate where increases in an interlinkage would decrease the leading eigenvalue λ_1 .




Fig. 6. Eigenvector sensitivities S^{eq} for (a) the ICSU matrix and (b) the GSDR matrix. Yellow and orange colours indicate that increasing those interlinkages serves to make the components of the leading eigenvector $\mathbf{v}^{(1)}$ more equal. Dark blue colours indicate that increasing an interlinkage will make the leading eigenvector more unbalanced. The relative differences between matrix entries are significant; the absolute differences between the colour scales in (a) and (b) are not significant.

internal network dynamics, then the internal dynamics of the network will re-orient itself towards more equal progress.

Taken together, Figs. 5 and 6 indicate that the two sensitivity measures, S^m and S^{eq} , defined above while having distinct interpretations, actually behave in very similar ways. As a result, there are specific perturbations to the interlinkage network that would both increase the growth rate λ_1 and equalise the leading eigenvector $\mathbf{v}^{(1)}$. For the ICSU interlinkage matrix, links from SDGs 1, 2 and 3 to SDGs 4, 5, 9 and 12 are optimal; for the GSDR matrix, links from SDGs 1, 2, 3 and 8 to SDGs 6 and 12 are optimal.

4. Integrated Assessment Models

A third recent example of an interaction matrix describing interlinkages between the SDGs is given by van Soest et al. (2019) in the context of Integrated Assessment Models (IAMs). The major focus of van Soest et al. (2019) is to review the construction and focus of IAMs, examine areas where they need to be further developed in order fully to account for all seventeen SDGs, and recommend how best to structure that future development.

Hence the paper focuses on IAMs as a set of tools rather than squarely focusing on the SDGs themselves.

As the authors remark early on in the paper, IAMs were originally developed primarily with a focus on climate, together with energy, economic variables, and land. Human development, for example as reflected in SDGs 5 (Gender Equality) and 11 (Inequality) tend to be far less well captured within IAMs in general. On this point of coverage of the SDGs, the authors found that four Goals SDGs 4 (Education), 5 (Gender Equality), 10 (Inequality), and 16 (Peace) were ‘clearly not well covered in these models’. SDG 17 on Partnerships is also difficult to incorporate within existing the methodology of existing IAMs.

The paper is underpinned by two surveys: a ‘model survey’ that focused on the question of the representation of SDGs within a set of 12 IAMs that participate in the Linking Climate and Development Policies – Leveraging International Networks and Knowledge Sharing (CD-LINKS) project, and an ‘expert survey’ to which 105 subject experts, including the authors of the ICSU report (International Council for Science (ICSU), 2015) and the GSDR 2019 (Independent Group of Scientists appointed by the Secretary-General, 2019), were invited. The expert survey asked

respondents to select one SDG that best covered their personal expertise, and then to indicate the linkages in both directions from this SDG to all the others, using the seven point scale proposed by Nilsson et al. (2016) and reproduced in Table 1 here for reference here.

From the 50 responses to the expert survey that were received, van Soest et al. (2019) construct an interaction matrix between the SDGs. With any data collection exercise there are always caveats and constraints that should be noted. For this expert survey there are four particular issues that should be noted that make it a less reliable set of interactions than the ICSU or GSDR networks presented above. First, compared to the ICSU and GSDR interaction matrices, this is a relatively small number of data points, given that each expert was asked to comment on linkages relating to a single SDG. Second, the expert responses were not subjected to further discussions, either in a Delphic method of allowing subsequent modification of scores through reaction to the views of other experts, or in a group discussion as was the case in the ICSU report. As a result it is difficult to verify that the experts applied common standards in their responses. Third, as van Soest et al. note, some respondents noted that individual targets within a Goal gave rise to synergies, and some to trade-offs, and so at the level of whole Goals the nature of the interactions were not clear enough to describe with confidence. Fourth, the supplementary material to van Soest et al. (2019) shows that the 50 expert opinions that were gathered were rather unequally distributed across the SDGs. For example, no experts focussed directly on SDG 5 (Gender Equality), and SDGs 7, 11 and 17 had the highest numbers of experts represented (six or seven in each case).

Finally, it is interesting to note that van Soest et al. included in their invitations all the authors of both the ICSU and GSDR 2019 reports. This indicates that as well as the potential variability in their results from working with a relatively small number of expert responses, there is some potential lack of independence in views between these expert responses and those that inform the ICSU and GSDR reports.

Mindful of these potential issues with the data from the expert survey, our methodology to construct an interlinkage network, starting from the data provided in the supplementary material to van Soest et al. (2019) was first, to re-weight the expert scores in order to adjust for the uneven response across the SDGs. This results in a more self-consistent set of average scores, where weights correspond to the proportions of available experts who gave positive and negative scores for different network links. As in the case of the GSDR analysis above, we consider separately the positive and negative scores in order to be able to run similar ensemble simulations to test the robustness of any conclusions and features of the network structure that arise. This could be extended to include the complete distribution of the weightings from -3 to $+3$ but this is unlikely to change significantly the conclusions of the analysis presented below.

We show in Fig. 7(a) an illustration of the adjacency matrix of the network, which we refer to as the IAM interlinkage network, in a form comparable to those in Fig. 1 for the ICSU and GSDR networks. There are interesting similarities in the overall structure, with a significant number of strongly positive (bright yellow) entries in the first three rows, showing the positive influences of all SDGs on SDGs 1, 2 and 3. The network indicates particularly large negative influences of SDG 9 (Industry) on SDGs 5 (Gender Equality), 14 (Life Below Water) and 15 (Life on Land). The negative influence of SDG 9 on SDG 5 results from a single expert view; only one expert respondent covered SDG 9 and (as remarked above) none covered SDG 5. So the negative influence $9 \rightarrow 5$ may not be fully representative of expert opinion more broadly, and indicates the need for care in drawing conclusions from the IAM matrix.

Fig. 7(b) does however provide a level of reassurance about the network structure overall since the leading eigenvalue of the IAM matrix $\lambda_1 = 6.975$ (3 d.p.) is much larger than the remaining eigenvalues; the next most positive is $\lambda_2 = 0.928$ (3 d.p.). This implies that the interactions will be dominated by those described by the leading eigenvector. This is also the case for the ICSU and GSDR interaction matrices but the separation of the leading eigenvalue from the rest is not so marked in those cases as it is for the IAM matrix.

This leading eigenvector is shown in Fig. 8(a) (the blue crosses) together with those for the ICSU matrix (red squares) and GSDR matrix (black dots). There are clear similarities in overall shape across the three: the first three components of the eigenvector are large, in fact for the IAM matrix they are the largest except for SDGs 4 (Education) and 9 (Industry), and the components corresponding to SDGs 14 and 15 are low, as they are also for the ICSU and GSDR eigenvectors. This indicates that the interlinkages described by the IAM matrix would lead to greatest progress on SDGs 1, 2, 4 and 9 and lower levels of progress on the other Goals. Fig. 8(b) examines the variation in the shape of the leading eigenvector under a similar (but not identical) process of randomization to that described in Section 2.4 for the GSDR matrix where both positive and negative values for links were available.

The difference between the simulations here and those described in Section 2.4 and Figs. 3 and 4 arises since there is a need here to explicitly take into account the zeros reported by experts in their assessments. The ensemble of 10^5 matrices A was generated by choosing entries of $+1$, -1 and 0 independently for each entry A_{ij}^{random} , for all entries where the number of expert opinions N_{ij}^{IAM} was positive. If $N_{ij}^{IAM} = 0$ then we set $A_{ij} = 0$; but if $N_{ij}^{IAM} > 0$ then we set A_{ij}^{random} to take the.

- value $+1$ with probability A_{ij}^+ ,
- value -1 with probability A_{ij}^- , and
- value 0 with probability $1 - A_{ij}^+ - A_{ij}^-$,

where A_{ij}^+ and A_{ij}^- are the weighted averages of the positive and negative expert scores, separately. We then compute the eigenvalues of A and the eigenvector $\mathbf{v}^{(1)}$ for the leading eigenvalue λ_1 . The corresponding violin plot is shown in Fig. 8(b). We see that the means of each eigenvector component (i.e. the horizontal black bars) agree with those plotted in Fig. 8(a) for the IAM matrix. The variabilities around these means are approximately equal across the 17 eigenvector components. As a result, a noticeable part of the distribution of the components for SDGs 14 and 15 is negative, indicating that these are most at risk when the whole ensemble of interaction matrices, and the separated effects of the positive and negative interlinkages, is considered.

The high values for the components of the leading eigenvector for SDGs 1 (No Poverty), 2 (Zero Hunger), 4 (Education) and 9 (Industry) also show up in the sensitivity matrices shown in Fig. 9 where the interlinkages that offer the greatest improvements are all links from one of those four SDGs. But it is interesting to note that the sensitivity S^m of the leading eigenvalue is highest when those linkages connect to SDGs 13 (Climate) or 16 (Peace), whereas for the sensitivity S^{eq} that attempts to equalise the components of the leading eigenvector the highest-scoring interlinkages are those that connect to SDGs 14 (Life Below Water) and 15 (Life on Land). One interpretation would be that there are more self-reinforcing loops containing SDGs 13 and 16, and that these are detected by S^m which is a measure only of the growth rate of the leading eigenvector. In contrast, S^{eq} is most improved by connections from the SDGs that have the highest eigenvector components to those that have the lowest components, as SDGs 14 and 15 do in this case.




Fig. 7. (a) Illustration of the IAM interlinkage network derived from the expert survey carried out in [van Soest et al. \(2019\)](#). Yellow and orange cells indicate high positive linkages, and dark blue cells indicate strongly negative interlinkages between each pair of SDGs $j \rightarrow i$. White cells indicate values close to zero. (b) Eigenvalues $\lambda_1, \dots, \lambda_n$ for the IAM interlinkage matrix (blue crosses) compared to those for the ICSU matrix (red squares) shown previously in [Fig. 2\(a\)](#).

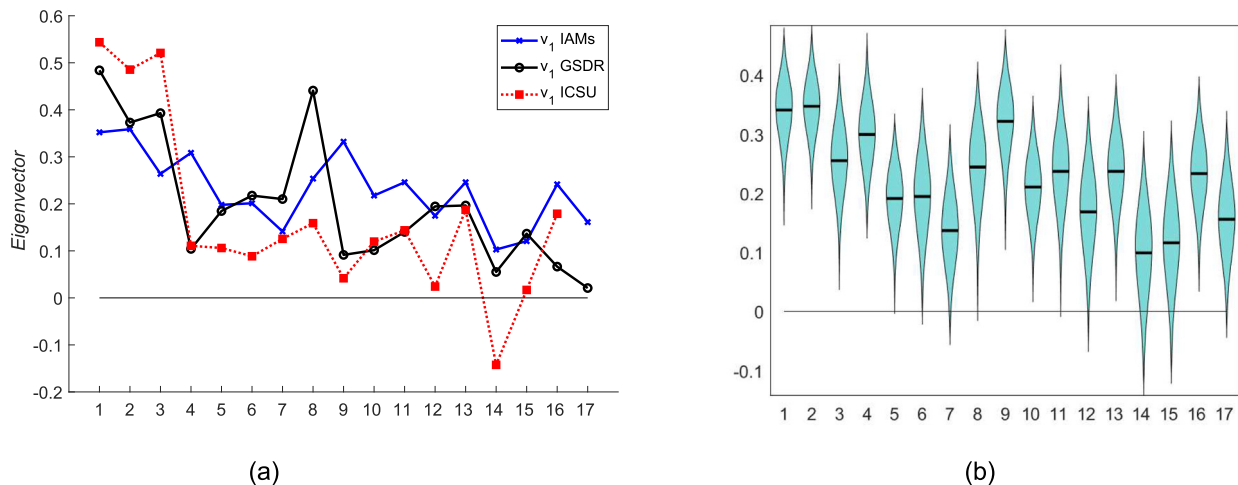


Fig. 8. (a) Leading eigenvectors for the ICSU, GSDR and IAM interlinkage matrices, each normalised so that $\sum_{i=1}^n (v_i^{(1)})^2 = 1$. Note the large components for SDGs 1, 2 and 3 in every case, and, likewise, the particularly low components on SDGs 14 and 15. (b) Violin plot for the IAM matrix illustrating the variability in the components $v_1^{(1)}, \dots, v_n^{(1)}$ of the leading eigenvector when the negative elements in the interlinkage matrix are taken into account. In both plots the horizontal axis labels the SDGs by their number.

5. Network hierarchy and overall directionality

In the discussion above we have repeatedly emphasised the directed nature of the networks of interlinkages. There is a significant difference between SDG j influencing SDG i and the converse. Mechanisms that enable influences in each the two directions are likely to be different, and so to relate to different policy contexts and priorities. This is particularly the case when the Goals related to 'human development', for example SDGs 5 (Gender Equality) and 10 (Inequality). Achieving Goals related to resource use (e.g. SDG 6 on Water, or SDG 7 on Energy) is possible in a variety of ways, of which many will not lead necessarily to progress on SDGs 5 and 10 as well. In this sense the interlinkages point to how likely it is that the mechanisms that are available to address one SDG are compatible with progress on others.

Capturing directionality in networks is a fundamental idea, and is most vividly illustrated in the literature on food webs in ecosystems, where nodes represent species, and directed edges describe which predator species prey on which others. In the ecosystems literature, reviewed recently by [Pringle and Hutchinson \(2020\)](#) key challenges include identifying precisely the trophic interactions themselves, as well as moving from the identification of individual predator-prey interactions to a system-wide analysis of the food web, along the way developing an idea of the structured layers within the food web. The expectation is often that such a stratification will make sense biologically, with layers containing similar types of prey that are preyed on by a collection of organisms in the next level up. Real food webs are driven by a wide range of external and internal effects, and the resulting behaviour is regulated in a complex manner involving species at all levels within




Fig. 9. Sensitivities for the Integrative Assessment Models (IAM) interaction matrix. (a) Eigenvalue sensitivity S^m . Yellow and orange colours indicate that increasing those interlinkages serves to increase the value of the leading eigenvalue. (b) Eigenvector sensitivity S^{eq} . Yellow and orange colours indicate network perturbations that will make the components of the leading eigenvector more equal.

the network (rather than just a few 'keystone' species) - a truly complex system (de Ruiter et al., 2005). The dynamics of real food webs are often highly dynamic, both in space and time; while here we appeal to food webs in order to motivate questions concerning the construction of a hierarchy within a directed network, we will restrict ourselves to a more static, but natural and general question: to what extent can a given directed network be organised into a collection of such 'trophic levels', so that the directed edges in the network point, as far as possible, all in the same direction. This question was recently explored by MacKay et al. (2020) and we follow their presentation in this section, applying their results to our SDG interlinkage networks.

Mathematically, the identification of network hierarchy, i.e. the relative levels of different nodes, can be set up as a minimisation problem. We wish to minimise the function

$$F(\mathbf{h}) = \frac{1}{\sum_{i,j=1}^n |A_{ij}|} \sum_{i,j=1}^n |A_{ij}| (h_i - h_j - 1)^2 \quad (5)$$

which depends on the vector of levels \mathbf{h} , where h_i is the layer height of node i , and the interlinkages A_{ij} . The theory presented by MacKay et al. (2020) assumes that the network has no negatively weighted edges, so we replace the interlinkage strength A_{ij} by its absolute value $|A_{ij}|$. The form of (5) indicates that $F(\mathbf{h})$ will be minimised by choices of the h_i that put a node i on a level (assumed to be spaced out roughly by the integers) below a node j , so that $h_i - h_j \approx 1$, when there is a directed edge $j \rightarrow i$. An explicit equation for the levels \mathbf{h} that minimise $F(\mathbf{h})$ can be deduced by differentiating (5) with respect to h_i and setting $\partial F / \partial h_i = 0$ for all i . This results in a linear matrix-vector equation which can be straightforwardly solved for the vector \mathbf{h} :

$$\Lambda \mathbf{h} = \mathbf{k}^{in} - \mathbf{k}^{out} \quad (6)$$

where we recall from Section 2.2 that $k_i^{in} := \sum_j A_{ij}$ is the in-degree of node i , $k_j^{out} := \sum_i A_{ij}$ is the out-degree of node j , and the Laplacian matrix Λ is defined to be $\Lambda := \text{diag}(\mathbf{k}^{in} + \mathbf{k}^{out}) - A - A^T$ where $\text{diag}(\mathbf{u})$ is the $n \times n$ matrix formed by putting the entries of the vector \mathbf{u} on the diagonal and zeros elsewhere.

For a given network, MacKay et al. (2020) refer to the minimum achievable value of $F(\mathbf{h})$ (denoted F_0) as its *trophic incoherence* and show that F_0 always lies between zero and 1. $F_0 = 0$ if and only if there exist a set of levels h_i that separate nodes into layers, spaced a unit apart, where all edges connect a node to (different) node in the layer above. At the other extreme, $F_0 = 1$ in the case that the graph is a union of equal-weight cycles, for which no separation into layers is possible and so the levels are equal, $h_1 = \dots = h_n$; this is maximal trophic incoherence and is consistent with there being no hierarchy that can be discerned among the nodes.

Fig. 10 shows the results of computing the optimal levels \mathbf{h} for each of the networks discussed above: the ICSU network, the GSDR 2019 network and the IAM network. Nodes that are lower down in the figures indicate SDGs that are further 'upstream' and in general influence other SDGs more than they themselves are influenced in return. Nodes at higher levels are 'downstream' and benefit from the influence of, and progress made on, other Goals. It should be noted that in the construction of these figures all interlinkages were taken to be positive, so that the distinction between trade-offs and co-benefits has been lost. However, the overwhelming majority of interlinkages are positive, and different treatments of the negative edges (for example omitting them completely) led to very similar results.

There are several general similarities between the three networks. First, SDGs 1 and 2 appear close to the top of each figure, suggesting that progress on these Goals benefits significantly from progress elsewhere. SDGs 12 and 17 are always close to the bottom of the figure (SDG 17 does not appear in Fig. 10(a) since it was not covered in the ICSU analysis). The very low position of SDG 7 (Energy) in the IAM network compared to its mid-ranking position in the ICSU and GSDR networks may well reflect the additional emphasis placed on SDG 7 in Integrated Assessment Models in general. As van Soest et al. (2019) noted, IAMs are generally well able to represent the influence of energy use on food production (SDG 2), water and sanitation (SDG 6), sustainable consumption (SDG 12), climate (SDG 13) and life on land (SDG 15). This viewpoint pushes attainment of SDG 7 (Energy) down towards the lowest levels of the network as its role 'enabling' other SDGs is emphasised by the construction and features of IAMs themselves.




Fig. 10. Trophic levels (on the y-axis) for the SDGs, derived from the three directed interlinkage networks: (a) ICSU network; (b) GSDR 2019 network; (c) IAM network. Low (high) trophic levels indicate Goals that put in (take out) more from the network than others. Red dots are labelled by the number of the SDG they represent; note that the ICSU network does not include SDG 17. Vertical coordinates in the plots are exactly the heights h_i for SDG i constructed in order to minimise $F(\mathbf{h})$, setting the lowest level to the value zero. Horizontal coordinates (x-axis) have no specific network interpretation and are chosen to optimise the clarity of the figures, for example by making as many edges as possible into straight lines (Brandes and Köpf, 2002). Edge thicknesses are proportional to the weights $|A_{ij}|$.

6. Discussion and conclusions

In this paper we have developed mathematical techniques to compute system-level features of interlinkage networks for the SDGs that could help to inform the development of more coherent policy responses for Agenda 2030. These mathematical techniques are perhaps the simplest possible that enable an integration of the individual-level description of the SDGs into a system-wide view. Their simplicity is both an asset and a liability; they allow overall conclusions to be extracted from individual links, but also reflect biases and unevenness when this is present in the construction of the underlying network. Nevertheless, they enable one to move beyond the mere construction of an interlinkage network towards the consideration of appropriate policy responses.

We also applied these techniques to compare three specific interlinkage networks for the SDGs. Each of these three is formed as a result of expert opinion in the policy or academic communities and so have different limitations to other classes of network, most obviously those based on correlations of time series of indicators for the SDGs.

In terms of results from the specific example networks we find that, broadly speaking, attention remains unequally distributed across the parts of Agenda 2030. The findings of the GSDR 2019 report show that some SDGs have clearly received significantly greater academic attention and policy scrutiny, with the result that the literature there is substantial while in other places it is considerably sparser, for example the literature on SDGs 4 (Education) and 5 (Gender equality). This is echoed in the context of Integrated Assessment Models where there is a clear emphasis on resource allocation. The result is, unsurprisingly, that the literature prioritises SDGs related to resources and the economy, for example SDGs 8 (Growth), 9 (Industry). In terms of network structure, SDGs 12 (Sustainable Consumption and Production) and 17 (Partnerships for the Goals) are consistently viewed as shaping the rest of Agenda 2030: these appear low down in the 'food webs' shown in Fig. 10, driving progress on the remaining SDGs. The reliance on SDG 7 (Energy) of the construction of Integrated Assessment Models is revealed by the much lower placing of SDG 7 in Fig. 10(c) compared to its relative position in Figs. 10(a) and (b).

All three networks have a dominant 'mode of response' that indicates that, without external inputs, the interlinkages among the SDGs that are proposed would drive the monotonic growth of progress on Agenda 2030 in one particular way, described by the leading eigenvector in each case. The alternative possibility, that the leading eigenvalues of the network are a complex conjugate pair, which would indicate oscillatory growth of a dominant mode with progress being naturally in-phase between some Goals and out of phase between others, does not arise although it is theoretically possible. This is reassuring; policies considered implicitly by the interlinkage networks are coherent at least in the most basic way possible in order to ensure the potential success of Agenda 2030. Indeed the eigenvalue plots in Figs. 2(a) and 7(b) indicate that the interlinkage matrices are (in some sense) reasonably far from admitting this oscillatory possibility. This also supports the stability that implicitly underpins data-driven analyses of SDG interactions (e.g. Pradhan et al., 2017) since it implies that trends should be monotonic and should persist over time.

The form of this response is given by the shape of the leading eigenvector, as shown in Fig. 2(b) and 8(a); for all three networks the components for SDGs 1, 2 and 3 are among the largest, indicating that the general directionality of the network is towards these three; other SDGs act to enable progress on SDGs 1–3 more than these three Goals are viewed as enabling progress elsewhere. Fig. 10 tells the same story, with SDGs 1–3 appearing towards the top of each plot, at the highest trophic levels, with other Goals pointing upwards towards them. We can conclude that the systemic views of the expert opinion reflected in these interlinkage networks are that SDGs 1–3 are promoted by the other Goals, asymmetrically. The linkages in the opposite directions, from SDGs 1–3 towards other Goals, are not strong enough to be able to conclude that Agenda 2030 is a single 'indivisible' collection of ambitions: it can be stratified, and the network analysis indicates how.

The sensitivity analysis of Section 3 indicates how addressing this stratification also improves either the growth rate of dominant mode of the network, or its equality along the different components (or both). Given the identification of the stratification of the SDGs into different levels (as shown in Fig. 10), the resolution makes intuitive sense: strengthen the couplings between progress

on SDGs 1–3 and other Goals (in particular SDG 12, but also SDGs 4, 5, 9, 14 and 17) in order to offset the stratification and integrate the network of SDGs with itself more comprehensively. This is the policy challenge.

In terms of risks, a recurring conclusion from the analysis, as shown in Figs. 2(b), 4, and 8(b) is that slow, or indeed negative, progress is most likely on the environmental Goals, in particular SDG 14 (Life below Water) but also SDG 15 (Life on Land). It is interesting to contrast the results for SDG 14 and SDG 13 (Climate). SDG 13 is far better supported by positive interlinkages, resulting in larger leading eigenvector components (Fig. 8a) and SDG 13 being placed at a higher level than SDG 14 in Fig. 10(a) and (b). The interlinkage networks systematically reinforce more progress on Climate than on Life below Water. The risk of a lack of progress on SDG 14 was highlighted by Dawes (2020) and it is noteworthy that it is an issue on which there is substantial agreement between the ICSU and GSDR networks.

Our final observation is on SDG 5 (Gender Equality): the ICSU report clearly views SDGs 4 (Education) and 5 as key inputs driving the SDG network, as is shown by the low position of these Goals in Fig. 10(a). However, the literature on SDG 5 as reviewed by the GSDR 2019 report appears to be considerably sparser than for other interlinkages. Tables 2 and 3 in Appendix A show that far fewer references to SDG 5 were made in the literature survey reported there than for other Goals. This may well indicate an important gap in our understanding of Agenda 2030 and that we do not have enough of an evidence base to understand how SDG 5 interacts with other Goals. This is also the case for the expert survey related to Integrated Assessment Models reported in van Soest et al. (2019) where no responses were obtained from experts who self-identified SDG 5 as their ‘field of expertise’. In the context of the GSDR 2019 report, if additional evidence were available which pointed to the positive impact of progress towards SDG 5 on other SDGs, including this would have the effect of moving SDG5 downwards in Fig. 10(b), so that its relative position became much more comparable those it occupies in Fig. 10(a) and (c).

Turning now to future work, we outline briefly three general important directions. First, the integration of data-driven approaches with the expert survey-driven interlinkage networks presented here. A key requirement for this is to isolate the directly self-reinforcing effects of progress on a goal in accelerating further progress on that goal, in order to separate this from the interlinkages between different SDGs. While the reports by the Sustainable Development Solutions Network (SDSN) (Sachs et al., 2019; Sachs et al., 2020) provide an excellent starting point, initial work by the author using the input data presented in the main reports, where levels and trend data are coarse-grained into one of only four states both for current levels of achievement of each SDG and the trend, has resulted in unreliable results containing large fluctuations. But the detailed and rich coverage of the SDSN report shows that further work could be extremely rewarding, allowing regional or country-level detection of important interlinkages that would help policymakers.

Indeed, many of the overall conclusions of the 2020 SDSN Report (Sachs et al., 2020) are echoed above. For OECD countries the SDSN report points to alarming trends on SDGs 13 and 14, and a lack of progress on SDG 5. Similar summary comments, particularly related to poor performance on SDGs 12–15, are made for most of the other regions of the world; see the summary comments in section 2.4 of Sachs et al. (2020).

More generally, there remains a significant opportunity to make use of the global indicator framework developed by the Inter-Agency and Expert Group on SDG Indicators (IAEG-SDGs) in the construction of interlinkage networks, for example incorporating the differences in uncertainty that are likely to exist between Tier 1 and Tier 2 indicators into robustness analyses. The use of SDG indicator data for the construction of interlinkages is not completely straightforward due to the need to estimate correlations in statistically appropriate ways, the need to distinguish between correlations and causal effects, the effects of data gaps (which are often very large) and instances in which even the best available indicator data series may not capture accurately the intention behind a particular target within Agenda 2030. We hope to return to this challenge in future work.

A second direction for future work is to use the sensitivity measures presented in Section 3 to inform the discussion of the likely effectiveness of the various proposed sets of policy actions that seek to condense Agenda 2030 into a more manageable framework. One example (among several) is the set of Six Transformations proposed by the SDSN. Each of the Six Transformations focuses on a subset of the SDGs and suggests, in essence, that the subset should be addressed together in policy terms. In terms of the interlinkage network this feels similar to increasing self-links from one SDG (or a small subset of SDGs) to themselves, in order to boost their progress. But this may miss more beneficial opportunities to link progress between SDGs that are currently not well-enough connected, for example designing healthcare provision in ways that automatically improve gender equality (i.e. seeking to increase the link from SDG 3 → SDG 5), which may well enable better system-wide progress over longer time horizons.

Thirdly, although we have investigated issues of robustness, most obviously in Section 2.4 with respect to the GSDR interaction matrix, there remain many other structural issues that merit further investigation. For example, are the conclusions reached from analysis at the level of whole goals also true for networks constructed from target-level data? To what extent is a lack of primary literature (‘evidence’) on a specific interlinkage actually evidence of the lack of that interlinkage, or could this be remedied by a more systematic approach to what appear, at first, to be missing data? How could we test different approaches to the construction of interlinkage networks for systematic overall biases towards a subset of the SDGs? While comparisons between different studies provide useful answers for some of these issues, it is clear that much remains to be done.

Declaration of Competing Interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

I am extremely grateful for fruitful discussions with, and encouragement from, Elizabeth Pollitzer and a number of colleagues especially Alastair Spence. I would also like to acknowledge and thank two anonymous reviewers whose insightful comments led to important improvements to the paper.

Appendix A

Table 2

Numbers N_{ij}^+ of positive interlinkages between Goals (i, j) reported in the GSDR 2019, where i is the row index and j the column index, and the (i, j) entry represents a positive influence of Goal j on Goal i . Row and column totals are shown in the final column and last row, respectively. Blanks indicate zero entries. Data shown also in [Pham-Truffert et al., 2020, Fig. 4.](#)

Goal j :	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	Total
Goal i																		
1	6	54	4	16	11	43	94	58	23	28	36	14	36	57	61	5	14	560
2	12	36	9	3	28	76	56	8	21	13	10	26	54	62	96	4	37	551
3	8	60	35	22	14	52	96	19	9	14	66	11	30	13	37	5	6	497
4	2	7	2	4	6	34	31	2	3	5	5		1	4		8	2	116
5	10	33	3	15	3	33	23	3	9	2	8	2	3	1	6	11	3	168
6		46	2		1	9	123	14	34	7	32	115	21	6	37	4	15	466
7		27	1	7		45	31	22	44	4	4	48	35	3	17	32	25	345
8	4	25	70	22	6	54	73	8	32	9	9	30	11	39	67	15	17	491
9		6		4		22	25		3	5	1	19	5	4	2	4	22	122
10		6	3	8	5	19	32	11	2	8	10	4	10	7	3	6	9	143
11				2		33	43		16	7	8	12	24	25	22	7	22	221
12		10		2	1	107	40	4	6	2	2	5	6	34	5	10	16	250
13		32	2	8		12	55	6	14	2	15	13	16	49	73	5	21	323
14				2		24	20	10	2	3	21	35	57	56	5	2	8	245
15		40	4	6	3	32	22	10	4	8	6	48	41	12	66	12	33	347
16	6	2		8	4	5	4	3	1	8	1	10	4	2	13	8	8	87
17				2		4		2	1	2	2	3	2		2	7	17	44
Total	48	384	135	131	82	604	768	180	224	127	236	395	356	374	512	145	275	4,976

Table 3

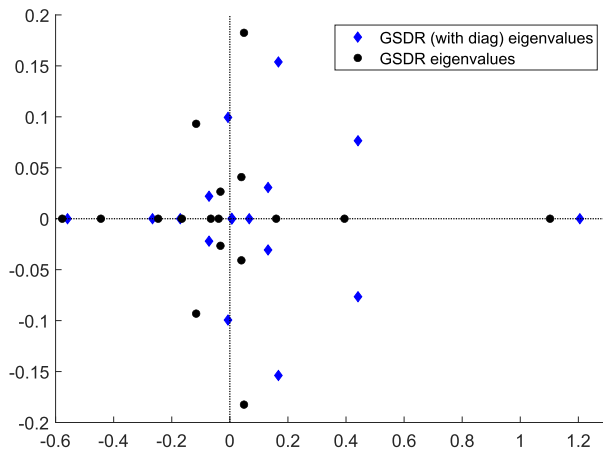
Numbers N_{ij}^- of reports of negative interlinkages between Goals (i, j) reported in the GSDR 2019, where i is the row index and j the column index, and the (i, j) entry represents a negative influence of Goal j on Goal i . Row and column totals are shown in the final column and last row, respectively. Blanks indicate zero entries. Data shown also in [Pham-Truffert et al., 2020, Fig. 4.](#)

Goal j :	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	Total
Goal i																		
1		7					30	3	2		4	4	1	11	9			71
2	1	8				1	28	1			1	2	3	6	3			54
3	4	25	1				19	2	1		2		1		1	3		59
4		4								1							1	6
5				1		1			1							2		5
6	19	50				2	61	10		6	4				4			156
7		15	1			9	2	3	4		2	1	3		27			67
8							21	2	2	2				14				41
9	1						1											2
10		1					16	4	3					2	1			27
11		2						1					1	19		1		24
12		3				1	3	4							1			12
13	1	11					7	8	1		2			2				32
14	2	1					9	9	3		22		75	7				128
15	10	39					26	1	4		9	1	2					92
16							3				1						2	6
17																		0
Total	38	166	2	1	0	14	226	48	21	9	47	8	86	61	46	6	3	782

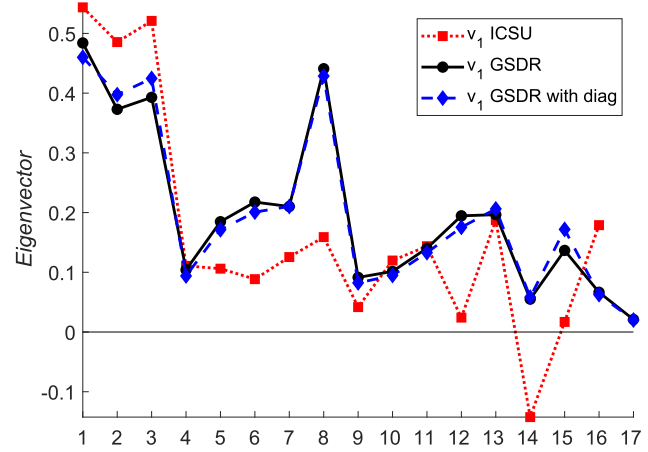
Table 4

Numbers $N_{ij}^{av} := N_{ij}^+ - N_{ij}^-$ (i.e. Table 3 subtracted from Table 2) of reports of interlinkages of either sign between Goals (i, j) reported in the GSDR 2019, where i is the row index and j the column index. Blanks indicate zero entries.

Goal j :	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
Goal i :																	
1	6	47	4	16	11	43	64	55	21	28	32	10	35	46	52	5	14
2	11	28	9	3	28	75	28	7	21	13	9	24	51	56	93	4	37
3	4	35	34	22	14	52	77	17	8	14	64	11	29	13	36	2	6
4	2	3	2	4	6	34	31	2	3	4	5		1	4		8	1
5	10	33	3	14	3	32	23	3	8	2	8	2	3	1	6	9	3
6	-19	-4	2		1	7	62	4	34	1	28	115	21	6	33	4	15
7		12		7		36	29	19	40	4	2	47	32	3	-10	32	25
8	4	25	70	22	6	54	52	6	30	7	9	30	11	25	67	15	17
9	-1	6		4		22	24		3	5	1	19	5	4	2	4	22
10		5	3	8	5	19	16	7	-1	8	10	4	10	5	2	6	9
11		-2		2		33	43	-1	16	7	8	12	23	6	22	6	22
12		7		2	1	106	37		6	2	2	5	6	34	4	10	16
13	-1	21	2	8		12	48	-2	13	2	13	13	16	47	73	5	21
14	-2	-1		2		24	11	1	-1	3	-1	35	-18	49	5	2	8
15	-10	1	4	6	3	32	-4	9		8	-3	47	39	12	66	12	33
16	6	2		8	4	5	1	3	1	8		10	4	2	13	8	6
17				2		4		2	1	2	2	3	2		2	7	17



(a)



(b)

Fig. 11. Comparisons between the interlinkage matrices from the GSDR 2019 Report (Independent Group of Scientists appointed by the Secretary-General, 2019) with and without the diagonal entries in the interlinkage matrix. (a) Eigenvalues of the (averaged) interlinkage matrices. (b) Components $v_1^{(1)}, \dots, v_n^{(1)}$ of the eigenvector $\mathbf{v}^{(1)}$ corresponding to the eigenvalue λ_1 having the largest real part, in the two cases, along with the leading eigenvector of the ICSU matrix for reference. The horizontal axis indicates the SDGs by number. Each eigenvector is normalised so that the root-mean-square of the entries is equal to one. We observe that the effect of the diagonal entries is small; for example the leading eigenvalue λ_1 increases from $\lambda_1 = 1.1026$ to $\lambda_1 = 1.2052$ (to 4 d.p.) when the diagonal entries are included.




Fig. 12. Histograms showing the probability distribution of the components of the leading eigenvector, for SDGs 1–9, for the ensemble of GSDR 2019 networks, illustrating the variability in the eigenvector components caused by the randomness in the choice of positive or negative entries in the matrix. These results correspond directly to the violin plot in Fig. 4.




Fig. 13. Histograms showing the probability distribution of the components of the leading eigenvector, for SDGs 10–17, for the ensemble of GSDR 2019 networks, illustrating the variability in the eigenvector components caused by the randomness in the choice of positive or negative entries in the matrix. These results correspond directly to the violin plot in Fig. 4.

Appendix B

In this Appendix we briefly set out results (omitting details and justifications), following [Greenbaum et al. \(2020\)](#), for the effects of perturbations on the eigenvalues and eigenvectors of a (square) matrix A . Let the matrix A have eigenvalues $\lambda_1, \dots, \lambda_n$, which we assume are distinct, with corresponding right and left eigenvectors $\hat{\mathbf{v}}^{(1)}, \dots, \hat{\mathbf{v}}^{(n)}$ and $\hat{\mathbf{y}}^{(1)}, \dots, \hat{\mathbf{y}}^{(n)}$ normalised so that $\hat{\mathbf{y}}^{(i)*}\hat{\mathbf{v}}^{(i)} = 1$ for all $i = 1, \dots, n$ where $*$ denotes the complex conjugate transpose. Without loss of generality we may also assume $\|\hat{\mathbf{v}}^{(i)}\| = 1$ for all i .

In the case that both sets of eigenvectors are normalised separately, we denote the right and left eigenvectors by $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(n)}$ and $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n)}$ where $\mathbf{y}^{(i)*}\mathbf{v}^{(i)} \neq 1$ in general, but $\|\mathbf{v}^{(i)}\| = \|\mathbf{y}^{(i)}\| = 1$ for all $i = 1, \dots, n$.

In order to compute the effect of a perturbation to adjacency matrix A we consider it to be a function of a parameter, writing $A(\varepsilon)$ where $A \equiv A(0)$. We suppose that $A(\varepsilon)$ has an eigenvalue $\lambda_1(\varepsilon)$ and a right eigenvector $\mathbf{v}^{(1)}(\varepsilon)$ that are continuously differentiable functions of ε , so that

$$A(\varepsilon)\mathbf{v}^{(1)}(\varepsilon) = \lambda_1(\varepsilon)\mathbf{v}^{(1)}(\varepsilon). \quad (7)$$

Differentiating with respect to ε and setting $\varepsilon = 0$ we obtain

$$A'(0)\mathbf{v}^{(1)}(0) + A(0)\mathbf{v}^{(1)'}(0) = \lambda_1'(0)\mathbf{v}^{(1)}(0) + \lambda_1(0)\mathbf{v}^{(1)'}(0), \quad (8)$$

where $'$ denotes a derivative with respect to ε . Now multiply on the left by the left eigenvector $\mathbf{y}^{(1)T}(0)$ and observe that the second term on the left hand side will cancel with the second term on the right hand side since they are both $\lambda_1(0)\mathbf{y}^{(1)T}(0)\mathbf{v}^{(1)'}(0)$. Hence we are left with

$$\lambda_1'(0) = \frac{\mathbf{y}^{(1)T}(0)A'(0)\mathbf{v}^{(1)}(0)}{\mathbf{y}^{(1)T}(0)\mathbf{v}^{(1)}(0)}. \quad (9)$$

In the special case that the perturbation is to a single network edge $j \rightarrow i$, the derivative matrix $A'(0)$ is just $A'(0) = \mathbf{e}_i\mathbf{e}_j^T$ (an outer product) where \mathbf{e}_i is the column vector whose i^{th} component is 1 and all other components are zero. For this case, omitting the parameter dependence in the notation, the expression (9) simplifies to become

$$\lambda_1' = \frac{y_i^{(1)} v_j^{(1)}}{\mathbf{y}^{(1)T}\mathbf{v}^{(1)}}$$

which describes the rate of change of the eigenvalue λ_1 with respect to the value of the network edge from $j \rightarrow i$. This result is well-known in the ecological literature, see for example Eq. (23) in [Neubert and Caswell \(1997\)](#) who derive it in the context of ecosystem resilience, i.e. considering the matrix A as describing perturbations from a stable equilibrium for which all eigenvalues $\lambda_1, \dots, \lambda_n$ have negative real parts, but it holds more generally as is evidenced by their citation of [Jacobi \(1846\)](#) as a historical reference to the result.

The matrix S^m introduced in (3) is precisely the $n \times n$ matrix of rates of change with respect to perturbations in each element of A in turn, scaled by a factor of λ_1 :

$$S_{ij}^m := \frac{1}{\lambda_1} \frac{\partial \lambda_1}{\partial A_{ij}} = \frac{1}{\lambda_1} \frac{y_i^{(1)} v_j^{(1)}}{\mathbf{y}^{(1)T}\mathbf{v}^{(1)}} \equiv \frac{1}{\lambda_1} \frac{\hat{y}_i^{(1)} \hat{v}_j^{(1)}}{\hat{\mathbf{y}}^{(1)T}\hat{\mathbf{v}}^{(1)}}. \quad (10)$$

The factor of $1/\lambda_1$ is included so that S^m computes relative changes in the magnitude of the leading eigenvalue: this allows better comparison of the values of S^m computed from different matrices with different absolute values of λ_1 . And the above calculation holds for both sets of eigenvectors: those with the 'hats' and

those without; as S_{ij}^m is a ratio of these terms, any rescaling of the eigenvectors in the numerator and denominator would just cancel out.

We turn now to consideration of perturbations to the eigenvectors. Let $V_1 = [\hat{\mathbf{v}}^{(2)} \dots \hat{\mathbf{v}}^{(n)}]$ and $Y_1 = [\hat{\mathbf{y}}^{(2)} \dots \hat{\mathbf{y}}^{(n)}]$ be the $n \times (n-1)$ matrices whose columns are the eigenvectors corresponding to $\lambda_2, \dots, \lambda_n$, normalised so that $\hat{\mathbf{y}}^{(i)*}\hat{\mathbf{v}}^{(i)} = 1$ in line with previous notation.

Further, let $D_1 = \text{diag}(\lambda_2, \dots, \lambda_n)$ be the $(n-1) \times (n-1)$ matrix whose diagonal elements are the remaining eigenvalues, and with all off-diagonal elements zero, and let I_{n-1} denote the $(n-1) \times (n-1)$ identity matrix, i.e. the matrix having ones on the main diagonal and zeros elsewhere. Greenbaum, Li & Overton note that the expression $V_1(D_1 - \lambda_1 I_{n-1})^{-1} Y_1^*$ is well-known and referred to by various names such as the *group inverse of $A - \lambda_1 I$* , or the *reduced resolvent matrix of A with respect to the eigenvalue λ_1* in different parts of the mathematical literature.

As before, let \mathbf{e}_i denote the vector that has i^{th} component equal to 1 and all other elements zero. Then, in the network cases considered here where the derivative of A with respect to this perturbation is just the matrix $\mathbf{e}_i\mathbf{e}_j^T$, we can use Eq. (7) in [Greenbaum et al. \(2020\)](#) to obtain a pleasingly compact expression for the eigenvector derivative as a linear combination of the other eigenvectors, with coefficients that can be written in terms of the eigenvalues and eigenvector components:

$$\begin{aligned} \hat{\mathbf{v}}^{(1)'} &\equiv \frac{\partial \hat{\mathbf{v}}^{(1)}}{\partial A_{ij}} = -V_1(D_1 - \lambda_1 I_{n-1})^{-1} Y_1^* \mathbf{e}_i \mathbf{e}_j^T \hat{\mathbf{v}}^{(1)} \\ &= \sum_{k=2}^n \hat{\mathbf{v}}^{(k)} \frac{\bar{y}_i^{(k)} \hat{v}_j^{(1)}}{\lambda_1 - \lambda_k}, \end{aligned} \quad (11)$$

where the first equality essentially follows from rearranging (8) and using the fact that

$$A(0) - \lambda_1 I_n = X \begin{bmatrix} 0 & 0 \\ 0 & D_1 - \lambda_1 I_{n-1} \end{bmatrix} Y^*$$

where $X \equiv [\hat{\mathbf{v}}^{(1)} V_1]$ and $Y \equiv [\hat{\mathbf{y}}^{(1)} Y_1]$; this can be verified directly. Details, including discussion of the underlying assumptions, are contained in [Greenbaum et al. \(2020\)](#). The second equality in (11) is just an algebraic simplification using the special form of \mathbf{e}_i and \mathbf{e}_j^T .

Since here we employ the standard normalisation for which both right and left eigenvectors are of unit length: $\|\mathbf{v}^{(i)}\| = 1$ and $\|\mathbf{y}^{(i)}\| = 1$ for $i = 1, \dots, n$ (this only leaves an ambiguity of the choice of sign in the case that eigenvalues are real), we now need to convert (11) into an expression for $\partial \mathbf{v}^{(1)}/\partial A_{ij}$, i.e. for the eigenvector $\mathbf{v}^{(1)}$ without the 'hat'.

This is also discussed in detail in [Greenbaum et al., 2020](#), see [Section 3.4](#), and we quote only the result here:

$$\frac{\partial \mathbf{v}^{(1)}}{\partial A_{ij}} = \frac{1}{(\hat{\mathbf{v}}^{(1)T}\hat{\mathbf{v}}^{(1)})^{1/2}} \frac{\partial \hat{\mathbf{v}}^{(1)}}{\partial A_{ij}} - \frac{(\partial \hat{\mathbf{v}}^{(1)}/\partial A_{ij})^T \hat{\mathbf{v}}^{(1)}}{(\hat{\mathbf{v}}^{(1)T}\hat{\mathbf{v}}^{(1)})^{3/2}} \hat{\mathbf{v}}^{(1)} \quad (12)$$

which, when combined with (11), shows that the eigenvector derivative is just a linear combination of the original eigenvectors. It is a rather complicated-looking linear combination in order to guarantee that the perturbed eigenvector is still normalised to have unit length, and therefore it can be directly compared with the original eigenvector. Now, finally, combining (12) with (4) enables the computation of S^{eq} starting from the eigenvalues λ_i and eigenvectors $\hat{\mathbf{v}}^{(i)}$ and $\hat{\mathbf{y}}^{(i)}$ of A .

References

- Allen, C., Metternicht, G., & Wiedmann, T. (2018). Initial progress in implementing the Sustainable Development Goals (SDGs): a review of evidence from countries. *Sustainability Science*, 13, 1453–1467.
- Bali Swain, R., & Ranganathan, S. (2021). Modeling interlinkages between sustainable development goals using network analysis. *World Development*, 138, 105136.
- Le Blanc, D., Freire, D.C., & Vierros, M. (2017). Mapping the linkages between oceans and other Sustainable Development Goals: A preliminary exploration. UN Department of Economic & Social Affairs, DESA Working Paper No. 149.
- Bonacich, P. (1972). Factoring and weighting approaches to clique identification. *Journal Mathematical Sociology*, 2, 113–120.
- Bonacich, P. (2007). Some unique properties of eigenvector centrality. *Social Networks*, 29, 555–564.
- Brandes, U., & Köpf, B. (2002). Fast and simple horizontal coordinate assignment. *Lecture Notes in Computer Science*, Volume 2265, 31–44. Springer.
- Breuer, A., Janetschek, H., & Malerba, D. (2019). Translating sustainable development goal (SDG) interdependencies into policy advice. *Sustainability*, 11, 2092.
- Dawes, J. H. P. (2020). Are the Sustainable Development Goals self-consistent and mutually achievable? *Sustainable Development*, 28, 101–117.
- de Ruiter, P.C., Wolters, V., Moore, J.C., & Winemiller, K.O. (2005). Food web ecology: playing Jenga and beyond. *Science* 309, 68–71 (2005).
- Greenbaum, A., Li, R.-C., & Overton, M. L. (2020). First-order perturbation theory for eigenvalues and eigenvectors. *SIAM Review*, 62(2), 463–482.
- Hoffmann, H. (2015). violin.m - Simple violin plot using MATLAB default kernel density estimation. INRES (University of Bonn), Katzenburgweg 5, 53115 Germany. Email: hhoffmann@uni-bonn.de. Retrieved from <https://uk.mathworks.com/matlabcentral/fileexchange/45134-violin-plot>.
- Independent Group of Scientists appointed by the Secretary-General. (2019). Global Sustainable Development Report 2019: The Future is Now - Science for Achieving Sustainable Development. United Nations, New York.
- International Council for Science (ICSU). (2017). A Guide to SDG Interactions: from Science to Implementation. Eds: D.J. Griggs, M. Nilsson, A. Stevance, D. McCollum. International Council for Science, Paris.
- International Council for Science (ICSU), International Social Sciences Council (ISSC). (2015) Review of the Sustainable Development Goals: The Science Perspective. Paris: International Council for Science (ICSU). ISBN: 978-0-930357-97-9.
- Jacobi, C. J. G. (1846). Über ein leichtes Verfahren die in der Theorie der Säcularstörungen vorkommenden Gleichungen numerisch aufzulösen. *J. Reine und Angewandte Mathematik*, 30, 51–95.
- Le Blanc, D. (2015). Towards integration at last? The sustainable development goals as a network of targets. *Sustainable Development*, 23, 176–187.
- Lusseau, D., & Mancini, F. (2019). Income-based variation in sustainable development goal interaction networks. *Nature Sustainability*, 2, 242–247.
- MacKay, R. S., Johnson, S., & Sansom, B. (2020). How directed is a directed network? *Royal Society Open Science*, 7, 201138.
- Nelson, R. B. (1976). Simplified calculation of eigenvector derivatives. *AIAA Journal*, 14(9), 1201–1205.
- Neubert, M. G., & Caswell, H. (1997). Alternatives to resilience for measuring the responses of ecological systems to perturbations. *Ecology*, 78(3), 653–665.
- Newman, M. (2018). *Networks: an introduction* (Second edition). Oxford University Press.
- Nilsson, M., Griggs, D., & Visbeck, M. (2016). Policy: map the interactions between sustainable development goals. *Nature*, 534, 320–322.
- Pedercini, M., Arquitt, S., Collste, D., & Herren, H. (2019). Harvesting synergy from sustainable development goal interactions. *PNAS*, 116, 23021–23028.
- Pham-Truffert, M., Rueff, H., & Messerli, P. (2019). Knowledge for Sustainable Development: Interactive repository of SDG interactions. CDEdatablog, Bern, Switzerland: CDE. Retrieved from <https://datablog.cde.unibe.ch/index.php/2019/08/29/sdg-interactions/>. Accessed on 7 September 2020.
- Pham-Truffert, M., Metz, F., Fischer, M., Rueff, H., & Messerli, P. (2020). Interactions among Sustainable Development Goals: Knowledge for identifying multipliers and virtuous cycles. *Sustainable Development*, 28, 1236–1250.
- Pradhan, P., Costa, L., Rybski, D., Lucht, W., & Kropp, J. P. (2017). A systematic study of Sustainable Development Goal (SDG) interactions. *Earth's Future*, 5, 1169–1179.
- Pringle, R. M., & Hutchinson, M. C. (2020). Resolving food-web structure. *Annual Review of Ecology Evolution and Systematics*, 51, 55–80.
- Ranganathan, S., & Bali Swain, R. (2018). Sustainable development and global emission targets: a dynamical systems approach to aid evidence-based policy making. *Sustainable Development*, 26, 812–821.
- Ranganathan, S., Nicolis, S. C., Bali Swain, R., & Sumpter, D. J. T. (2017). Setting development goals using stochastic dynamical system models. *PLoS ONE*, 12(2), e0171560.
- Raworth, K. (2017). *Doughnut economics: seven ways to think like a 21st-century economist*. Random House.
- Sachs, J., Schmidt-Traub, G., Kroll, C., Lafortune, G., & Fuller, G. (2019). Sustainable Development Report 2019: G20 and Large Countries edition. New York: Bertelsmann Stiftung and Sustainable Development Solutions Network (SDSN).
- Sachs, J., Schmidt-Traub, G., Kroll, C., Lafortune, G., Fuller, G., & Woelm, F. (2020). *The Sustainable Development Goals and COVID-19. Sustainable Development Report 2020*. Cambridge: Cambridge University Press.
- Schaeffer, P.V., & Willardsen, K. (2019). A Note on the Tinbergen Rule. Working paper. Available from <https://www.petervschaeffer.com/publications.html>.
- Singh, G. G., Cisneros-Montemayor, A. M., Swartz, W., Cheung, W., Guy, J. A., Kenny, T., McOwen, C. J., Asch, R., Geffert, J. L., Wabnitz, C. C. C., Sumaila, R., Hanich, Q., & Ota, Y. (2018). A rapid assessment of co-benefits and trade-offs among Sustainable Development Goals. *Marine Policy*, 93, 223–231.
- Spaiser, V., Ranganathan, S., Bali Swain, R., & Sumpter, D. J. T. (2016). The sustainable development oxymoron: quantifying and modelling the incompatibility of sustainable development goals. *International Journal of Sustainable Development & World Ecology*, 24, 457–470.
- Tinbergen, J. (1952). *On the Theory of Economic Policy*. New York: North-Holland.
- United Nations (2015). *Transforming our world: The 2030 Agenda for Sustainable Development A/RES/70/1*. New York: United Nations General Assembly.
- van Noordwijk, M., Duguma, L. A., Dewi, S., Leimona, B., Catacutan, D. C., Lusiana, B., Öborn, I., Hairiah, K., & Minang, P. A. (2018). SDG synergy between agriculture and forestry in the food, energy, water and income nexus: reinventing agroforestry? *Current Opinion in Environmental Sustainability*, 34, 33–42.
- van Soest, H. L., van Vuuren, D. P., Hilaire, J., Minx, J. C., Harmsen, M. J. H. M., Krey, V., Popp, A., Riahi, K., & Luderer, G. (2019). Analysing interactions among Sustainable Development Goals with Integrated Assessment Models. *Global Transitions*, 1, 210–225.
- Weitz, N., Carlsen, H., Nilsson, M., & Skånberg, K. (2018). Towards systemic and contextual priority setting for implementing the 2030 Agenda. *Sustainability Science*, 13, 531–548.