Environ Syst Decis (2014) 34:49–59
DOI 10.1007/s10669-014-9489-2

Research prioritization using hypothesis maps

Alan Masinter · Mitchell Small · Elizabeth Casman

Published online: 19 February 2014

Abstract This work presents a method to aid in the prioritization of research within a scientific domain. The domain is encoded into a directed network in which nodes represent factors in the domain, and directed links between nodes represent known or hypothesized causal relationships between the factors. Each link is associated with a numeric weight that indicates the degree of understanding of that hypothesis. Increased understanding of hypotheses is represented by higher weights on links in the network. Research is prioritized by calculating optimal allocations of limited research resources across all links in the network that maximize the degree of overall knowledge of the research domain. We quantify the level of knowledge of individual nodes (factors) in the map by a network centrality measure that reflects in dependencies between knowledge level of nodes and the knowledge level of their parent nodes in the map. We analyzed a funded research proposal concerning the fate and transport of nanomaterials in the environment to illustrate the method.

Keywords Research prioritization · Nanomaterial fate and transport · Causal maps

1 Introduction

Setting priorities for research is a difficult and recurring exercise for scientists and scientific funding agencies. Individuals and organizations involved in the funding and administration of scientific activities must decide which investments have the greatest potential to achieve the objectives of the research endeavor. However, for research domains with highly interrelated and poorly understood components, it is difficult to quantify expected future benefits from research. In such cases, priorities are typically determined through a variety of qualitative or deliberative approaches. This work presents a tool to aid in the evaluation of research priorities that capitalizes on interrelationships and dependencies between different parts of the research. We encode the hypotheses of the research domain into a directed network and posit a mechanism by which knowledge is transferred through the network as a function of research on specific hypotheses in the network. Such hypothesis maps explicitly show how knowledge of one factor is dependent on knowledge of other factors, and how those factors are in turn dependent on knowledge of other factors, and so forth, potentially continuing indefinitely. This recursive structure can be exploited to determine the best allocation of new research effort in order to facilitate the flow of information through the graph. Information can "flow" through the graph when it is not obstructed at any particular link, and the structure of the map will determine which links are most critical for information flow. By optimizing the flow of information through the hypothesis map for incrementally increasing levels of resource investment, we prioritize the links in the map for study.

Electronic supplementary material The online version of this article (doi:10.1007/s10669-014-9489-2) contains supplementary material, which is available to authorized users.

A. Masinter · M. Small · E. Casman (✉)
Department of Engineering and Public Policy, Carnegie Mellon University, 129 Baker Hall, 5000 Forbes Ave, Pittsburgh, PA 15213, USA
e-mail: casman@andrew.cmu.edu

2 Background

2.1 Scientific research assessment and prioritization

Research prioritization is the selection of research projects based on any number of quantitative and qualitative criteria. Quantitative prioritization approaches involve some attempt to measure the expected benefits, costs, and risks associated with investment decisions, and to select those projects which are expected to have the highest overall value, either singly or as a portfolio (Information System Laboratories 2003).

Prioritization with quantitative methods alone can bias outcomes toward research with only measurable, tangible benefits (Fischhoff 2000), and these methods can have issues associated with their reliability, validity, cost, timeliness, and acceptability (Feller and Stern 2007). Quantitative prioritization methods are particularly problematic for complex systems where it is difficult to quantify expected outcomes. Prioritizing scientific research is difficult in part because the goals of scientific funding organizations typically involve furthering knowledge and understanding of some scientific domain, rather than more quantifiable outputs (e.g., sales, patents). Additionally, a portfolio of research projects is likely to involve projects with interrelated outcomes, in that the results from one project are likely to alter what can be learned from other projects.

Some quantitative approaches have been applied to scientific research, though these instances involve some ultimate impact on a quantifiable objective beyond scientific understanding alone. For example, value-of-information approaches have been applied to research regarding environmental health risk management (Yokota and Thompson 2004) and to assessments of which of research investments would be most useful for achieving environmental goals [e.g., (Canis et al. 2010)].

Prioritization systems that are more qualitative might include informal discussion and review, voting systems, or science and technology road-mapping (Kostoff and Schaller 2001). However, deliberative approaches can be biased both by the selection of participants and by the framing of the questions on research priorities that are posed to them (Feller and Stern 2007).

Bibliometric methods can also be used to assess research progress by tracking and analyzing citations between published papers or patents. Combining bibliometric analysis of citation networks with the keywords associated with each article can identify important trends in research content (Kajikawa et al. 2010). In some fields that are particularly amenable to representation as a network (e.g., protein–protein interaction [Ganapathiraju and Orii 2013]), bibliometric methods have been used to inform future research priorities. However, the majority of bibliometric applications have been limited to retrospective analyses, and are of limited utility for planning future research.

This work presents a method for determining priority research topics, but successful implementation of a research prioritization method requires attention to factors outside the research material alone, such as the bureaucratic structure of the research organization, the management of equipment and personnel, or the mechanisms of communication and collaboration between researchers. Scientific organizations may be concerned with generating data that would support new proposals, supporting particular individuals' endeavors, or other considerations. In this work, we ignore the details of execution that are particular to any one organization, focusing only on the determination of the most effective areas for advancing knowledge.

2.2 Cognitive, causal, and hypothesis maps

Complex phenomena are commonly diagrammed as sets of linked influences or factors. Causative chains of factors form a directed network, and models using such networks of factors include all forms of influence diagrams, expert systems, mental modeling, cognitive mapping, and causal mapping. All of these involve variables or other separable components connected by directed links in a network. For example, Morgan constructed an influence diagram composed of causal relationships between variables related to nanomaterial properties and risks (Morgan 2005).

The term *concept mapping* in the evaluation and planning literature refers to a distinct approach for identifying concepts within a domain and depicting their relationships (Trochim 1989). In this approach, many participants group concepts together, and these results are combined to form a single, most representative grouping. Results can also be used to construct undirected networks of concepts for further analysis (McLinden 2013).

A cognitive map is a directed network that represents a person's assertions about some limited conceptual domain. The nodes of cognitive maps denote qualitative or quantitative variables, and the links between nodes are typically associated with a sign denoting positive or negative influence (Axelrod 1976; Bougon et al. 1977; Eden and Ackermann 1992; Ford and Hegarty 1984). Fuzzy cognitive maps have variable weights on links between nodes that denote the degree of positive or negative influence between elements (Kosko 1986) and have been used to gain insight into many phenomena, including ecological systems (Özesmi and Özesmi 2004; Prigent et al. 2008). Causal maps are similar to cognitive maps, but only contain links that represent statements of causality (Narayanan and Armstrong 2005). We use the term *hypothesis map* to describe a variant of a causal map in which links between nodes represent known or hypothesized influences between factors in a research domain. A link from one node to another indicates that there is at least one known or hypothesized way in which the first factor influences the second. All types of links are allowed in hypothesis maps; this includes cycles, such as when a node links to itself or when two nodes link to each other.

2.3 Construction of causal maps

Causal maps can be constructed from one or more text sources that describe the domain of interest in detail. These text sources may be transcripts of interviews conducted for the purpose of causal mapping, or may be documents written for another purpose. To construct a causal map, all causal assertions made within one or two sentences in the document are identified, including both explicit and implicit statements of causality. The causal statements are encoded into nodes and links between nodes, corresponding to concepts and concept relationships, respectively, while maintaining the original language of the source as faithfully as possible (Axelrod 1976; Narayanan and Armstrong 2005). The resulting causal map will often contain concepts and relationships that are redundant due to synonyms or closely related terms used in multiple causal statements. To simplify the map, related concepts are merged into groups. This reduces the number of nodes and may reduce the number of links in the network. Nelson et al. found that three levels of granularity are useful for this process (Narayanan and Armstrong 2005): (a) *concepts*, the original words and phrases from the source document; (b) *categories*, groups of multiple related concepts, joining synonyms or other terms that are closely related; and (c) *constructs*, groups of related categories. Grouping concepts leads to simpler and more interpretable maps at the cost of some detail. The grouping of concepts must be informed by knowledge of the research domain in order to preserve meaning of the nodes and relationships between them while simplifying the map. One approach for increasing the fidelity of the map is to have multiple researchers independently encode causal statements repeatedly until sufficient concordance is reached (Narayanan and Armstrong 2005). Statements of causality found in descriptive texts commonly assert relationships between concepts that are only indirectly related, and all such statements must be disentangled to construct an accurate hypothesis map. Additionally, some detail is necessarily lost when converting the full text description of the research into a hypothesis map, both through the initial encoding of the text into a network of linked concepts, and through any further grouping of concepts. Automated word clustering approaches may be useful if sufficiently voluminous texts describing the research are available, but manual clustering may be necessary to preserve scientifically meaningful concept groups.

3 Methods

3.1 Construction of hypothesis maps describing environmental fate and transport of nanomaterials

We constructed a hypothesis map using a single collaborative research proposal for the Center for the Environmental Implications of Nanotechnology (CEINT), written in 2008. CEINT is a consortium of investigators from Duke University, Carnegie Mellon University, Howard University, Virginia Tech, the University of Kentucky, and Stanford University, sponsored by the National Science Foundation and US Environmental Protection Agency.

One of the major CEINT research themes is the study of nanomaterial transport and transformations in the environment. The primary tasks of this research theme are to (a) connect the chemistry and size of nanoparticles to their aggregation state and movement in the environment, (b) characterize biological and chemical transformations of such materials, and (c) use that understanding of nanomaterial transformations to predict their environmental persistence, transport, and bioavailability. The section of the proposal that described the scientific thinking on nanomaterial fate and transport was used to construct our illustrative hypothesis map.

For our purposes, it is not important whether or not the hypotheses in this proposal represent the "truth" about nanomaterial fate and transport, but only that the hypotheses accurately represent the understanding of nanomaterial transport as of the time the proposal was written. Decisions concerning research allocations must necessarily be based on current understanding of the field at the time the research prioritization decision is being made, and so any priority research topics suggested by this analysis would be valid in 2008, but would change as knowledge resulting from subsequent research accumulated. A hypothesis map of this type cannot be validated, nor does it need to be. It is a representation of hypotheses, not fact.

The CEINT research proposal was analyzed for all statements of causality related to the fate and transport of nanomaterials in the environment. A causal relationship between two factors was assumed if statements in the proposal indicated that one factor influenced another. Explicit and implicit statements of known or hypothesized causal relationships were identified by the phrases "will likely impact," "is the basis for understanding," "plays a critical role in determining," and other similar statements of influence. The causal relationships from this process were graphed as a network of factors with the links between nodes in the network representing the known or hypothesized causal relationships between factors.

The concept-level causal map was also converted to a category-level and construct-level hypothesis maps with the procedure for grouping nodes described in Sect. 2.3.

The current degree of understanding of each link in the construct-level hypothesis map was given a numerical value obtained through expert elicitation. We conducted three interviews, each with a different co-author of the CEINT proposal, asking them for subjective opinions concerning the state of knowledge, and importance of individual links of the research map.

3.2 Information flow and centrality measures for analysis of hypothesis maps

Mathematically, the hypothesis map is a graph $G(V, E)$ where V is a set of n vertices or nodes corresponding to factors in a research domain, and E is a set of m edges consisting of ordered pairs of nodes that represent known or hypothesized causal relationships from parent nodes to child nodes. The network may have any structure, including cycles; for example, the networks derived from the CEINT proposal had instances of reciprocal causality between pairs of nodes (e.g., transformations of nanoparticles influencing their aggregation and vice versa) and of nodes influencing themselves (e.g., some type of nanoparticle attachment and aggregation influencing another aspect of attachment and aggregation).

The graph is equivalently described by an $n \times n$ adjacency matrix A whose (i, j) entry a_{ij} is 1 if there exists a directed link from node i to node j, and is 0 if there is not.

The knowledge of any given node is partially dependent on the information from its upstream parent nodes in the network. Knowledge of the parent nodes (node knowledge scores) combined with knowledge of the mechanism of that influence (link knowledge ratings) contributes to knowledge of the child node. Let node knowledge scores be defined as $x_1, \ldots, x_n$ and link knowledge ratings as l_{ij} for pairs of directed links (i, j) in the network. Link knowledge ratings represent the level of understanding of the relationship between parent and child nodes. Link knowledge ratings range from 0 (no knowledge) to 1 (complete knowledge). The goal of research prioritization is to optimally increase the information that propagates through the network by conducting research on, and increasing the knowledge rating of each link.

Node knowledge scores x_i are written as a function of parent node scores and incoming link ratings.

We assume that the contribution from a given incoming node and link is the product of the incoming link score and corresponding node score ($x_i \times l_{ij}$). That is, both upstream knowledge and the mechanism of influence must be understood for the incoming link to contribute to downstream knowledge.

Second, for lack of evidence to the contrary, we assume multiple incoming node-link pairs contribute independently to downstream knowledge. That is, node scores are a function of the sum of the incoming links as in Eq. 1 (rather than some other function for combining contributions from incoming links).

$$x_j = f\left(\sum_{i=1}^{n} l_{ij} x_j\right) \quad j = 1 \ldots n \tag{1}$$

If the function f consists of multiplication by a proportionality constant, then the set of Eq. 1 can be written in matrix form as Eq. 2, where the components of W (the weighted adjacency matrix) at (i, j) are equal to $a_{ij} \times l_{ij}$, and W' is the transpose of W.

$$\mathbf{x} = \alpha W' \mathbf{x} \tag{2}$$

If the matrix W contained only values of 0 and 1, then this expression would be equivalent to the formula for calculating eigenvector centrality, where the proportionality constant α equals the inverse of the largest eigenvalue of W'. Eigenvector centrality is one way to compute the importance of each node in a graph. The eigenvector centrality of each node is proportional to the sum of the centralities of the nodes which are connected to it. The more central a node's parents, the more central is the node itself.

However, eigenvector centrality and similar measures using weighted links are problematic for networks with parentless nodes, since the links from such nodes can never be valued. The parentless nodes can have no knowledge to contribute to downstream nodes (Bonacich and Lloyd 2001). To avoid this problem, we assign all nodes in the network an exogenous *a priori* level of knowledge defined by a vector $\mathbf{e}$. Thus, nodes without incoming links have their exogenous *a priori* level of knowledge to contribute to their child nodes. In this paper, we set $\mathbf{e}$ to a vector of ones, implying equal *a priori* knowledge for each node, though other initial knowledge scores could be used, if known.

Therefore, our final formulation of the node knowledge score is defined by Eq. 3.

$$\mathbf{x} = \alpha W' \mathbf{x} + \mathbf{e} \tag{3}$$

This is equal to the expression for alpha-centrality (Bonacich and Lloyd 2001) except that the adjacency matrix A is replaced by a weighted matrix W (Newman 2004).

An exact solution for $\mathbf{x}$ in Eq. 3 is described by Eq. 4, but involves matrix inversion, which can be computationally burdensome for large matrices.

$$\mathbf{x} = (I - \alpha W')^{-1} \mathbf{e} \tag{4}$$

However, if W has a distinctly largest eigenvalue, then, for values of $|\alpha| < \lambda^{-1}$, Eq. 4 has an equivalent infinite sum representation can be computed to within acceptable accuracy relatively quickly (Golub and van Van Loan 1983) (Eq. 5).

$$\mathbf{x} = \left(I + \sum_{t=1}^{\infty} \alpha^t W^t\right)' \mathbf{e} \tag{5}$$

In Eq. 5, t is the "walk length" of directed walks in the graph ending in node x. (A directed walk is a connected sequence of nodes and links, oriented in the same direction, without restrictions on revisiting nodes or links).

This formulation of weighted alpha-centrality scores shows how this measure takes into account the structure of the network as a whole while discounting distant connections more than close ones. Contributions from walks of every possible length are added together in the infinite sum across values of t, and when $|\alpha| < 1$, longer walks through the network contribute less than shorter walks by a factor of α for each additional step.

Walks in a hypothesis map represent the causal linkages between concepts and all their distant and near contributing factors in terms of the knowledge embodied in the linkages at a given time. Changing the knowledge of any link changes the knowledge flow through all nodes downstream of that link. The network structure will influence which changes have the largest impact.

For sufficiently large values of $|\alpha|$, the infinite sum in Eq. 5 is not convergent. In particular, the absolute value of α must be smaller than the inverse of the largest eigenvalue (λ) of W. We set α equal to the inverse of the largest eigenvalue of A, such that node knowledge scores are generally defined and convergent unless research on the map is no longer possible, i.e., when all $l_{ij} = 1$.

3.3 Prioritization of hypothesis map links for investment of resources

Increasing link knowledge ratings represents increasing understanding of the hypotheses that the links represent. Here, we assume that the objective of new research is to maximize the total knowledge of all factors. Thus, we calculate sets of link knowledge ratings that maximize the sum of node knowledge scores. Other objectives, such as maximizing selected node scores or maximizing the minimum node score, are conceivable. We felt that maximizing the sum of the node knowledge scores reflected the researchers' objective of understanding the entire field of nanomaterial fate and transport. The minimum value for a node knowledge score is 1, for a node without any contribution from parent nodes to its score (from the vector $\mathbf{e}$ in Eq. 3). The maximum value for a node knowledge score is unbounded; the score will diverge to infinity as all of the node's upstream link knowledge ratings approach 1. In order to normalize the contribution of each node knowledge score, each score is adjusted with a scaling function $S(x) = (x - 1)/x$, which normalizes the range of node knowledge scores from $[1, \infty)$ to the range $[0, 1]$.

To calculate optimal allocations of resources, link knowledge ratings are chosen to maximize the sum of scaled node knowledge scores. Let $\mathbf{L} = \{l_{ij,k} : (i,j) \in E\}$ be the set of m link knowledge ratings assigned to edges in the hypothesis map. A sequence of optimal allocations is calculated, indexed by k, by choosing values for link knowledge ratings $l_{ij,k} \in \mathbf{L}$ under different resource constraints R_k.

There are two constraints on these optimizations: first, link knowledge ratings cannot fall below their *a priori* values ($l_{ij,0}$), nor can they rise above 1; second, the total cost of increasing link knowledge ratings $[C(l_{ij,k})]$ cannot exceed the resource constraint R_k. Equation 6 describes the problem.

$$\begin{aligned} &\underset{l_{ij,k}\in\mathbf{L}}{\text{maximize}} && \sum_{i=1}^{n} \frac{x_i - 1}{x_i} \\ &\text{subject to} && l_{ij,0} \le l_{ij,k} \le 1 \quad \forall (i,j) \in E, \\ & && \sum_{l_{ij,k}\in\mathbf{L}} C(l_{ij,k}) \le R_k \end{aligned} \tag{6}$$

3.3.1 Cost function and initial conditions

We set the incremental cost of increasing link knowledge ratings to be constant and equal across links by using the cost function $C(l_{ij,k}) = l_{ij,k}$, though other cost functions could be used, if known. For simplicity, we set all *a priori* link knowledge ratings ($l_{ij,0}$) to 0. We found that specifying different values for *a priori* link knowledge ratings had little effect on overall results (see the Supplementary Information).

3.4 Prioritizing hypotheses for study to maximize total system knowledge

Optimal research allocations were calculated using numerical derivative-based local maximization over all link knowledge ratings ($\mathbf{L}$), incrementing R_k under each iteration. This results in a series of optimal allocations for different levels of funding that can be used to prioritize hypotheses. Initial values of $l_{i,j,k}$ in the sequential optimizations were set to the previous optimal value ($l_{i,j,k-1}$). We found no evidence of multiple local maxima for this network.

The priority of a given link is calculated from the series of optimal allocations by calculating a *link priority score* for each link, calculated by Eq. 7.

$$s_{ij} = \sum_{k=1}^{m} \left(1 - \frac{R_k}{R_m}\right) \left(l_{ij,k} - l_{ij,k-1}\right) \tag{7}$$

This measure of link priority calculates the differences in link knowledge ratings between sequential optimizations (second term) and weights these differences by how early they occur (first term). The lowest possible link priority score is 0, if only the final mth optimization (highest investment level) moves l_{ij} from 0 to 1. Its maximum possible value is close to 1 $(1 - R_1/R_m)$ if the first optimization yields $l_{ij,1} = 1$. (In this case, since *a priori* link knowledge ratings are 0, then $l_{ij,0} = 0$, and the second term equals 1). Otherwise, the amounts of investments are weighted by the total amount of resources available when they occur. When a link's priority score is close to 1, investments in the link are made despite low values of R_k (i.e., a small budget) and the link is high priority. If a link's priority score is close to 0, investments in the link are made only at high values of R and is thus low priority.

3.4.1 *Prioritization of research related to the fate and transport of manufactured nanomaterials in the environment*

In this work, we illustrate the use of hypothesis maps by applying it to the study of the fate and transport of manufactured nanomaterials in the environment. Manufactured nanomaterials are increasingly used in a wide array of consumer products and industrial processes, and pose largely unquantified risks to the environment. The uncertainties regarding nanomaterial risks and the rapid proliferation of nanomaterials have made the study of nanomaterial risk an important and urgent issue (Klaine et al. 2008). To understand nanomaterial risks, it is critical to understand the physical, chemical, and biological processes that may transport and transform nanomaterials in the environment, as these processes determine the concentration and form of nanomaterials to which biota are exposed and form the basis of quantitative risk assessment.

Since 2007, more than 30 publications have identified research gaps, outlined research strategies, or otherwise contributed to efforts related to the prioritization of research related to environmental, health, and safety aspects of nanomaterials (NRC 2012). One of the more comprehensive nanomaterial research strategy documents is the National Nanotechnology Initiative (NNI) report (NSET/NEHI 2011), which identified funding gaps and details broad areas of research required for comprehensive environmental risk assessment and risk management of nanomaterials. Another overview of nanomaterial research needs was a report published in 2012 by the National Academy of Sciences (NAS) (NRC 2012) on funding gaps and priorities, which highlighted the need to determine how particular nanomaterial properties affect key processes occurring in the environment, the standardization of nanomaterials and nanomaterial assays, the choice of which nanomaterials to study, modeling efforts, data collection, and other strategic research issues. The United States Environmental Protection Agency has also attempted to prioritize research questions regarding several specific nanomaterial applications through a series of workshops in which experts used a sequential voting procedure to rank critical research questions (Davis et al. 2010).

3.5 Hypothesis map for the fate and transport of nanomaterials

We identified 182 statements of causal relations (links) among 172 unique concepts related to the fate and transport of nanomaterials from the CEINT research proposal. The concept network contained 20 disconnected subgraphs, including 16 subgraphs containing only two or three nodes. By collapsing synonymous concept nodes into *category* nodes, the graph was reduced to 46 nodes and 131 links, with all but 4 nodes (2 links) within a single connected graph. Further grouping of categories into *constructs* resulted in a single connected construct network consisting of 10 nodes with 40 links. Figure 1 shows a comparison of the concept, category, and construct maps side-by-side to illustrate the consequences of collapsing nodes.

4 Results

4.1 Research allocation results

We ran sequential optimizations for research allocations for both the construct-level and category-level networks. In the category network, results indicate that it is most important to understand how nanomaterial transformations interact with aggregate structure (ranks 1, 2, 5, and 7). Links connected to the "aggregate structure" node (excepting one) comprised the top 20 priority links, followed by links with "nanoparticle properties" as parent node (ranks 21 through 26). Low priority links were typically related to the physical movement of nanoparticles. The 10 highest priority links, along with examples of corresponding quotes from the research proposal, are shown in Table 2. A summary of the category links, sorted by the quartile of their priority scores, is given in the Supplementary Information.

In the construct network, the 11 highest priority links in all involve the "nanoparticle properties and transformations" node. This node has high centrality, as it affects 7 other nodes in the network and is affected by itself and 3 other nodes out of the 10 total nodes in the network. The

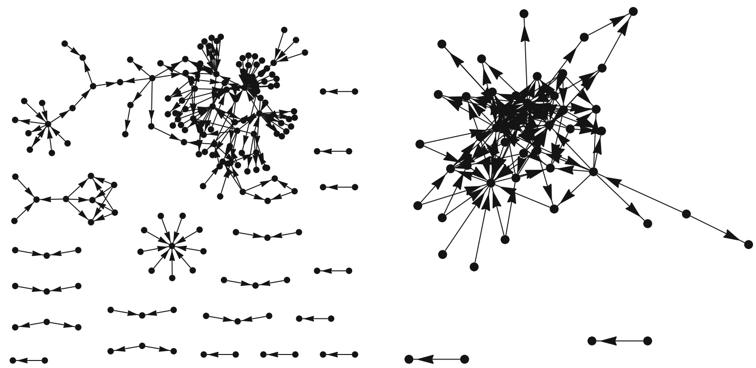

(a) Concept network. **(b)** Category network. **(c)** Construct network.

Fig. 1 Three levels of hypothesis maps constructed from the proposal section on the fate and transport of nanomaterials. The proposal contained 182 statements of dependencies between 172 distinct concepts, illustrated in the concept-level graph (A). The category-level graph (B) contained 131 influences between 46 nodes after grouping similar concepts into categories. The construct-level graph (C) results from a further grouping of categories and contains 40 links between 10 distinct constructs

Fig. 2 Link priority scores (LPS) for the construct-level hypothesis map. Low is defined as $0 \leq LPS \leq 1/3$, Medium as $1/3 < LPS \leq 2/3$, and High as $2/3 < LPS \leq 1$

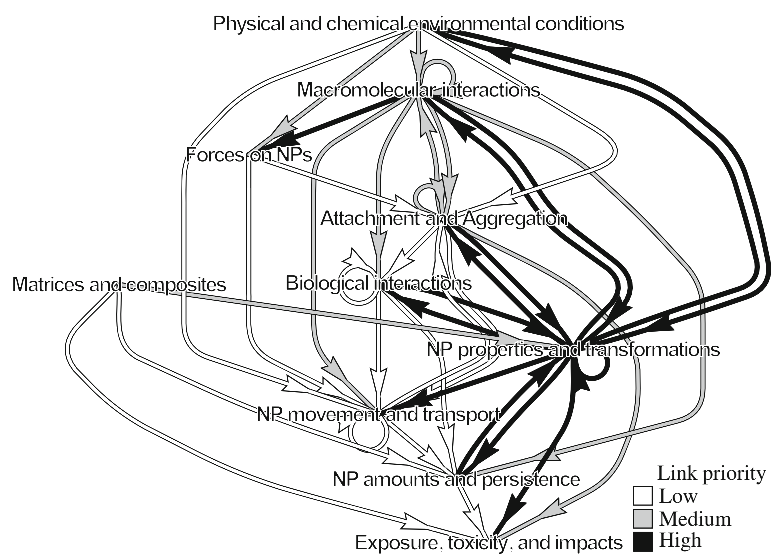

links with the lowest priority scores involved effects of the "matrices and composites" node on its child nodes. A diagram of the construct network, with links colored by priority score, is shown in Fig. 2. The 10 highest priority links, along with examples of corresponding quotes, are listed in Table 3. The links, sorted by the quartile of their priority scores, are presented in Table 1. *Optimization results for the construct-level map show that understanding nanomaterial transformations in the environment is most critical for understanding the fate of transport of nanomaterials, and ultimately for understanding the risks posed by nanomaterials in the environment*. Results from both the construct and category networks indicate the need to fully understand the fundamental chemical properties of nanoparticles and their transformations in the environment, including their aggregation behavior, in order to understand the other parts of the research domain.

4.2 Alternative optimization formulations

In calculating optimal research allocations, we used cost functions such that the cost of incrementing link knowledge ratings was constant and equal across links. Optimal allocations calculated under this constraint are indicative of the value of increasing knowledge in each area, but not necessarily the cost-effectiveness of investing there. Other cost functions, if known, could be used to more realistically simulate research investments.

5 Discussion

Six years after the original proposal was written, after much of the work was completed, we asked the original experts for their reactions to the optimization we performed on their

Table 1 Summary of construct-level links and their priorities

Parent node		Quartile of priority score			
No.	Name	4th (high priority)	3rd	2nd	1st (low priority)
1	Attachment and aggregation	–	5, 9	1, 2, 3, 8	–
2	Biological interactions	–	9	–	2, 7, 8
3	Exposure, toxicity, and impacts	–	–	–	–
4	Forces on NPs	–	–	1	8
5	Macromolecular interactions	9	1, 4, 5, 7	2, 8	–
6	Matrices and composites	–	9	–	3, 7
7	NP amounts and persistence	–	9	3	–
8	NP movement and transport	–	–	1	7, 8
9	NP properties and transformations	1, 2, 3, 5, 7, 8, 9, 10	–	–	–
10	Physical and chemical environmental conditions	9	4	5	1, 8

Each node in the construct-level map is listed in the first column, along with an assigned number. Child nodes of this node are listed by number in the other columns, divided into quartiles by their link priority score

Table 2 The ten highest-priority links of the category-level hypothesis map, sorted by average link priority score

Parent node	Child node	Example quote	Priority score
NP properties	Aggregate structure	The primary tasks in Theme 1 are to … connect the chemistry… of NPs to their aggregation state	0.99
NP reactivity	Aggregate structure	… we hypothesize that reactivity of nanomaterials is intimately related to aggregate structure	0.97
Transport	Aggregate structure	Aggregate structure in turn depends on those processes leading to aggregation, in particular, the transport processes	0.95
Aggregation	Aggregate structure	However, if aggregation leads to the formation of porous aggregates…	0.94
Aggregate structure	NP properties	More dense aggregates of NPs have a shadowing effect of neighboring particles	0.93
Aggregate structure	Macromolecules	… explore the role of NP aggregate morphology on interactions with well-defined macromolecules	0.92
Aggregate structure	NP reactivity	More dense aggregates of NPs … are hypothesized to have lower reactivities	0.92
Aggregate structure	Aggregation Kinetics	Equations 2–4 show particle aggregation rate to be a strong function of aggregate structure	0.91
Aggregate structure	Transport	Dependence of… transport… on the structure of aggregates they form	0.91
Aggregate structure	Aggregation	More dense aggregates of NPs have reduced mass transfer into the aggregate	0.91

2008 proposal. Although in 2010 they were not in agreement regarding the most important hypotheses, by 2014 they agreed that nanomaterial transformations were most central for understanding the fate and transport of nanomaterials in the environment. To quote one expert, "if you don't get transformations right, you can't get anything else right." Numerous papers resulting from the grant support this statement (Auffan et al. 2010; Auffan et al. 2009; Chang and Vikesland 2009; Cheng et al. 2011; Dale et al. 2013; Erbs et al. 2010; French et al. 2013; Gondikas et al. 2012; Kaegi et al. 2011; Kent and Vikesland 2012; Kirschling et al. 2011; Labille et al. 2010; Levard et al. 2012; Levard et al. 2013; Levard et al. 2013; Levard et al. 2011; Li et al. 2010; Louie et al. 2013; Lowry et al. 2012; Lowry et al. 2012; Ma et al. 2013; Phenrat et al. 2008; Rebodos and Vikesland 2010; Reinsch et al. 2012; Reinsch et al. 2010; Tiwari and Marr 2010; Wiesner et al. 2011; Wirth et al. 2012). This is not a formal validation, but it was gratifying to note that the prioritization results were supported by the experts after the completion of the research.

The low priority given to links influencing the "exposure, toxicity, and impacts" node, arguably the most

Table 3 The ten highest-priority 10 links of the construct-level hypothesis map, sorted by link priority score

Parent node	Child node	Example quote	Priority score
NP properties and transformations	NP properties and transformations	Transformations of nanomaterials in the environment may affect NP surface chemistry and therefore NP chemical reactivity	0.96
Macromolecular interactions	NP properties and transformations	We hypothesize that adsorbed macromolecules will reduce the innate reactivity of NPs (such as ROS production)…	0.88
NP properties and transformations	Macromolecular interactions	The ratio of trains to tails and loops (i.e. adsorbed layer conformation) will depend on… the nanomaterial properties.	0.87
NP properties and transformations	Physical and chemical environmental conditions	Transformations of nanomaterials in the environment may affect NP surface chemistry and therefore NP chemical reactivity… Conversely, nanomaterials and other chemical species they interact with may be altered as a result of these interactions.	0.84
NP properties and transformations	Attachment and aggregation	we hypothesize that reactivity of nanomaterials is intimately related to aggregate structure	0.84
NP properties and transformations	Biological interactions	Interactions between biological macromolecules and NPs receive special attention as we hypothesize… that NPs may modify macromolecular properties (such as DNA expression and protein configuration)	0.83
NP properties and transformations	NP amounts and persistence	Transformations of nanomaterials in the environment may affect NP surface chemistry and therefore… persistence	0.83
NP properties and transformations	Exposure, toxicity, and impacts	Photolysis, oxidation, and subsequent dissolution (in the case of mineral NPs) can: release toxic metals into the environment	0.83
NP properties and transformations	NP movement and transport	Modifications to NP surfaces, engineered or unintentional, will likely alter exposure and hazard through impacts on mobility	0.83
Environmental conditions	NP properties and transformations	… quantify the redox properties of individual NPs and determine environmental variables (e.g., pH and redox state) that enhance or decrease reactivity	0.8

important node of the hypothesis map, deserves some explanation. This result is an artifact of focusing on the proposal's section on nanomaterial fate and transport research. Another set of research topics in the CEINT research proposal is focused on nanomaterial exposure routes, biological responses and toxicity, and ecological impacts. That section has its own complex hypothesis map "downstream" of the fate and transport map, but is represented as a single node ("exposure, toxicity, and impacts") that influences nothing else in the fate and transport construct map. As a result, increasing its knowledge score contributes relatively little to the overall goal of understanding the fate and transport of nanomaterials in the environment. Its low score was an artifact of truncating the hypothesis map at that node.

The structure of the network is expected to change as new hypotheses are made or disproved, and new links are added or removed from the hypothesis map. In practice, it would be prudent to periodically re-evaluate established priorities as new information is generated. We assumed we could model how link knowledge ratings might change with additional investments, but research outcomes are highly uncertain and research does not always succeed. Additional research might even reveal the area to be more complex than previously believed, in effect lowering link knowledge ratings in the short term. In this work, we set the costs of incrementing link knowledge ratings to be constant and equal across links, such that results were representative of the value of additional knowledge, but we do not account for the cost of obtaining it or the likelihood of success. The incorporation of elicited cost functions for each link would be one way to deal with this problem, but it is difficult for anyone, experts included, to quantitatively predict future research outcomes. We found that results were fairly insensitive to variations in the formulation of the objective function; however, there is no guarantee that this will be the case for other hypothesis maps (see the Supplementary Information).

Hypothesis mapping can show what areas are important to understand from the point of information flow, but this will not always be the dominant consideration in prioritization decisions. We elicited subjective research priorities from three domain experts in 2010 and found our results to be uncorrelated with theirs. This suggests that the mapping method may contribute a novel perspective when used in conjunction with standard deliberative methods of research prioritization. Because of the complexity of research allocation decisions, we do not envision that results from hypothesis mapping would be the sole factor in such decisions. Rather, hypothesis mapping might be used in combination with other prioritization approaches.

In contrast to the example that we have provided, in which we prioritized a research endeavor where the hypothesis map was relatively rich, funding agencies may be interested in cutting edge science, where hypothesis maps are more abstract and less detailed. We expect that it would be difficult to obtain hypothesis maps that represent the current state of the science in such a field. If, however, a trusted initial hypothesis map could be generated and resources allocated, initial research results could then be used to update the hypothesis map for re-optimization. Iterative solutions are the hallmark of optimization, after all, and research prioritization is no exception.

Acknowledgements Support for this research was provided by the Department of Education under Grant Award P200A090055-11, the National Science Foundation and the Environmental Protection Agency NSF Cooperative Agreement EF-0830093, the Center for the Environmental Implications of NanoTechnology (CEINT), and Carnegie Mellon University. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the ED, NSF or the EPA. This work has not been subjected to EPA review, and no official endorsement should be inferred.

References

Auffan M, Pedeutour M , Rose J , Masion A, Ziarelli F, Borschneck D, Chaneac C, Botta C, Chaurand P, Labille J et al. (2010) Structural degradation at the surface of a tio2-based nanomaterial used in cosmetics. Environ Sci Technol 44(7):2689–2694

Auffan M, Rose J, Wiesner M, Bottero J-Y (2009) Chemical stability of metallic nanoparticles: a parameter controlling their potential cellular toxicity in vitro. Environ Pollut 157(4):1127–1133

Axelrod RM (1976) Structure of decision: the cognitive maps of political elites. Princeton University Press, Princeton

Bonacich P, Lloyd P (2001) Eigenvector-like measures of centrality for asymmetric relations. Soc Netw 23(3):191–201

Bougon M, Weick K, Binkhorst D (1977) Cognition in organizations: an analysis of the Utrecht Jazz Orchestra. Adm Sci Q 22(4): 606–639

Canis L, Linkov I, Seager TP (2010) Application of stochastic multiattribute analysis to assessment of single walled carbon nanotube synthesis processes. Environ Sci Technol 44(22):8704–11

Chang X, Vikesland PJ (2009) Effects of carboxylic acids on nc_{60} aggregate formation. Environ Pollut 157(4):1072–1080

Cheng Y, Yin L, Lin S, Wiesner M, Bernhardt E, Liu J (2011) Toxicity reduction of polymer-stabilized silver nanoparticles by sunlight. J Phys Chem C 115(11):4425–4432

Dale AL, Lowry GV, Casman EA (2013) Modeling nanosilver transformations in freshwater sediments. Environ Sci Technol 47(22):12920–12928

Davis JM, Wang A, Shtakin JA (2010) Nanomaterial case studies: nanoscale titanium dioxide in water treatment and in topical sunscreen. US EPA, Research Triangle Park, NC

Eden C, Ackermann F (1992) The analysis of cause maps. J Manage Stud 29(3):309–324

Erbs JJ, Berquó TS, Reinsch BC, Lowry GV, Banerjee SK, Lee PR (2010) Reductive dissolution of arsenic-bearing ferrihydrite. Geochim Cosmochim Acta 74(12):3382–3395

Feller I, Stern PC (eds) (2007) A strategy for assessing science: behavioral and social research on aging. The National Academies Press, Washington, DC

Fischhoff B (2000) Scientific management of science? Policy Sci 33(1):73–87

Ford JD, Hegarty WH (1984) Decision makers' beliefs about the causes and effects of structure: an exploratory study. Acad Manag J 27(2):271–291

French RA, Monsegue N, Murayama M, Hochella Jr. MF (2013) The structure and transformation of the nanomineral schwertmannite: a synthetic analog representative of field samples. Phys Chem Miner 1–10

Ganapathiraju MK, Orii N (2013) Research prioritization through prediction of future impact on biomedical science: a position paper on inference-analytics. GigaScience 2(11)

Golub GH, van Van Loan CF (1983) Matrix computations. John Hopkins University Press, Baltimore

Gondikas AP, Morris A, Reinsch BC, Marinakos SM, Lowry GV, Hsu-Kim H (2012) Cysteine-induced modifications of zero-valent silver nanomaterials: implications for particle surface chemistry, aggregation, dissolution, and silver speciation. Environ Sci Technol 46(13):7037–7045

Information System Laboratories (2003) Formal methods of decision analysis applied to prioritization of research and other topics, nureg/cr-6833. Technical report

Kaegi R, Voegelin A, Sinnet B, Zuleeg S, Hagendorfer H, Burkhardt M, Siegrist H (2011) Behavior of metallic silver nanoparticles in a pilot wastewater treatment plant. Environ Sci Technol 45(9):3902–3908

Kajikawa Y, Takeda Y, Matsushima K (2010) Computer-assisted roadmapping: a case study in energy research. Foresight 12(2): 4–15

Kent RD, Vikesland PJ (2012) Controlled evaluation of silver nanoparticle dissolution using atomic force microscopy. Environ Sci Technol 46(13):6977–6984

Kirschling TL, Golas PL, Unrine JM, Matyjaszewski K, Gregory KB, Lowry GV, Robert DT (2011) Microbial bioavailability of covalently bound polymer coatings on model engineered nanomaterials. Environ Science Technol 45(12):5253–5259

Klaine SJ, Alvarez PJJ, Batley GE, Fernandes TF, Handy RD, Lyon DY, Mahendra S, McLaughlin MJ, Lead JR (2008) Nanomaterials in the environment: behavior, fate, bioavailability, and effects. Environ Toxicol Chem 27(9):1825–1851

Kosko B (1986) Fuzzy cognitive maps. Int J Man Mach Stud 24(1): 65–75

Kostoff RN, Schaller RR (2001) Science and technology roadmaps. IEEE Trans Eng Manage 48(2):132–143

Labille J, Feng J, Botta C, Borschneck D, Sammut M, Cabie M, Auffan M, Rose J, Bottero J-Y (2010) Aging of tio< sub> 2</sub> nanocomposites used in sunscreen. dispersion and fate of the degradation products in aqueous environment. Environ Pollut 158(12):3482–3489

Levard C, Matt HE, Lowry GV, Brown Jr. GE (2012) Environmental transformations of silver nanoparticles: impact on stability and toxicity. Environ Sci Technol. 46(13):6900–6914

Levard C, Hotze EM, Colman BP, Dale AL, Truong L, Yang XY, Bone AJ, Brown Jr. GE, Tanguay RL, Di Giulio RT et al. (2013) Sulfidation of silver nanoparticles: natural antidote to their toxicity. Environ Sci Technol 47(23):13440–13448

Levard C, Mitra S, Yang T, Jew A, Badireddy AR, Lowry GV, Brown GE (2013) The effect of chloride on the dissolution rate of silver

nanoparticles and toxicity to *E. coli*. Environ Sci Technol 47(11):5738–5745
Levard C, Reinsch BC, Marc MF, Oumahi C, Lowry GV, Brown Jr. GE (2011) Sulfidation processes of pvp-coated silver nanoparticles in aqueous solution: impact on dissolution rate. Environ Sci Technol 45(12):5260–5266
Li Z, Greden K, Alvarez P, Gregory KB, Lowry GV (2010) Adsorbed polymer and nom limits adhesion and toxicity of nano scale zerovalent iron to *E. coli*. Environ Sci Technol 44(9):3462–3467
Louie SM, Tilton RD, Lowry GV (2013) Effects of molecular weight distribution and chemical properties of natural organic matter on gold nanoparticle aggregation. Environ Sci Technol 47(9): 4245–4254
Lowry GV, Espinasse BP, Badireddy AR, Richardson CJ, Reinsch BC, Bryant LD, Bone AJ, Deonarine A, Chae S, Therezien M et al. (2012) Long-term transformation and fate of manufactured ag nanoparticles in a simulated large scale freshwater emergent wetland. Environ Sci Technol 46(13):7027–7036
Lowry GV, Gregory KB, Apte SC, Lead JR (2012) Transformations of nanomaterials in the environment. Environ Sci Technol 46(13):6893–6899
Ma R, Levard C, Marc MF, Brown Jr. GE , Lowry GV (2013) Sulfidation mechanism for zinc oxide nanoparticles and the effect of sulfidation on their solubility. Environ Sci Technol 47(6):2527–2534
McLinden D (2013) Concept maps as network data: analysis of a concept map using the methods of social network analysis. Eval Program Plann 36(1):40–8
Morgan K (2005) Development of a preliminary framework for informing the risk analysis and risk management of nanoparticles. Risk analysis: an official publication. Soc Risk Anal 25(6):1621–35
Narayanan VK, Armstrong DJ (2005) Causal mapping for research in information technology. IGI Global, Hershey, PA
National Research Council (NRC) (2012) Committee to develop a research strategy for environmental, health, and safety aspects of engineered nanomaterials. A research strategy for environmental, health, and safety aspects of engineered nanomaterials
Newman MEJ (2004) Analysis of weighted networks. Phys Rev E 70(5):56131
NSET/NEHI (2011) NNI Environmental health, and safety research strategy. Technical report, National Science and Technology Council Committee on Technology, Subcommittee on Nanoscale Science, Engineering, and Technology
Özesmi U, Özesmi SL (2004) Ecological models based on people's knowledge: a multi-step fuzzy cognitive mapping approach. Ecol Model 176(1–2):43–64
Phenrat T, Long TC, Lowry GV, Veronesi B (2008) Partial oxidation (aging) and surface modification decrease the toxicity of nanosized zerovalent iron. Environ Sci Technol 43(1):195–200
Prigent M, Fontenelle G, Rochet M, Trenkel VM (2008) Using cognitive maps to investigate fishers' ecosystem objectives and knowledge. Ocean Coast Manag 51(6):450–462
Rebodos RL, Vikesland PJ (2010) Effects of oxidation on the magnetization of nanoparticulate magnetite. Langmuir 26(22): 16745–16753
Reinsch BC, Levard C, Li Z, Ma R, Wise A, Gregory KB, Brown Jr. GE , Lowry GV (2012) Sulfidation of silver nanoparticles decreases escherichia coli growth inhibition. Environ Sci Technol 46(13):6992–7000
Reinsch BC, Forsberg B, Lee PR, Kim CS, Lowry GV (2010) Chemical transformations during aging of zerovalent iron nanoparticles in the presence of common groundwater dissolved constituents. Environ Sci Technol 44(9):3455–3461
Tiwari AJ, Marr LC (2010) The role of atmospheric transformations in determining environmental impacts of carbonaceous nanoparticles. J Environ Qual 39(6):1883–1895
Trochim WMK (1989) An introduction to concept mapping for planning and evaluation. Eval Program Plann 12:1–16
Wiesner MR, Lowry GV, Casman E, Bertsch PM, Matson CW, Di GRT, Liu J, Hochella Jr. MF (2011) Meditations on the ubiquity and mutability of nano-sized materials in the environment. ACS Nano 5(11):8466–8470
Wirth SM, Lowry GV, Tilton RD (2012) Natural organic matter alters biofilm tolerance to silver nanoparticles and dissolved silver. Environ Sci Technol 46(22):12687–12696
Yokota F, Thompson KM (2004) Value of information analysis in environmental health risk management decisions: past, present, and future. Risk analysis: an official publication. Soc Risk Anal 24(3):635–650