

Collaborative Creation of Teaching-Learning Sequences and an Atlas of Knowledge

Nagarjuna G. Gnowledge Laboratory, HBCSE, TIFR, India.

nagarjun@gnowledge.org

ABSTRACT: The article is about a new online resource, a collaborative portal for teachers, which publishes a network of prerequisites for teaching/learning any concept or an activity. A simple and effective method of collaboratively constructing teaching-learning sequences is presented. The special emergent properties of the dependency network and their didactic and epistemic implications are pointed. The article ends with an appeal to the global teaching community to contribute prerequisites of any subject to complete the global roadmap for an atlas being built on similar lines as Wikipedia. The portal is launched and waiting for community participation at <http://www.gnowledge.org>.

1. Meaning is in the links.

We have undertaken to study the problem of meaningful learning by focussing on the structure of knowledge as a semantic network.¹ A semantic network is a declarative graphic notation for representing knowledge in patterns of interconnected nodes and links [19]. Our epistemological standpoint is: Meaning of any concept that we harbour or any activity that we do as cognitive agents is generated out of links between other concepts or activities; Correct understanding is considered a mark of proper linkages between nodes, and misunderstanding is a mark of improper linkages; No node in such a network is meaningful in its own right, but only by virtue of the links the node has with the neighbouring nodes. This *neighbourhood theory of meaning* is the point of departure for our analysis, methodology, and the nature of the teaching-learning resources that we produce.² Our focus is on the relations between the nodes, because

1 Similar undertakings on meaningful learning, inspired by Ausubel [2], gave rise to Novak's concept mapping methodology [14], currently widely used in the science and mathematics education community.

2 Apart from the intellectual influences mentioned above, the neighbourhood theory of meaning is also a philosophical culmination of the naturalism of Quine [7]. A fuller statement of the theory is unpublished and under active development by the author.

they provide clues to understanding as well as misunderstanding.

Several kinds of semantic networks are distinguished and used in the literature spread across disciplines. Currently our focus in the article shall be on a simple methodology of generating teaching-learning sequences using the semantic network technique, followed by the emergent properties of such a network and their implications for the teaching-learning process (didactics) with marginal notes on epistemological implications.

The main purpose of this short communication is to share the excitement of the already evident results and inform the community about the on going work, that we hope will have serious implications to teaching and learning practices. Significance of the didactic implications suggested in the article depends a lot on the comprehensiveness of the knowledge base, which cannot be achieved without the participation of the global teaching community. Thus, this article serves the purpose of inviting the community with the promise of fruitful results.

2. Teaching-learning sequences using the dependency relation.

The dominant learner-centric cognitive constructivist stand point in education recommends to start from what the student already knows, namely the *prior knowledge*. Several researchers in the recent past have also turned our attention to the influence of prior-knowledge as the cause of misconceptions, as well as a requirement for constructing new knowledge [6], [20]. Based on our epistemological assumption mentioned previously, that understanding as well as misunderstanding arises due to correct or incorrect relations (reconstruction) between nodes, it is important to look at the nature of the relations. Of the many relations we focus on only one of them, *dependency relation*, to construct a directed graph which represents the flow of meaning. With the assumption that the flow of meaning cannot be different from the path of learning, we carry on the project of building the dependency maps, which can also be called as teaching-learning sequences.

2.1 The dependency relation generates a road map of all knowledge.

Every good teacher ensures that the student has the prerequisite *prior knowledge* before introducing any new topic. Most of the good text books explicitly mention the *prerequisites* at the beginning of each

chapter. This has been the generic guiding principle of curriculum design. This principle also stands out as one of the consensual principles of student centered teaching practices. We start from this consensus and attempt to consolidate it so as to make it a firm foundation on which rests all engagements of knowledge.

Let us gather the prerequisites and make a semantic network containing a single relation, *dependency*. If P is a prerequisite of Q , then we consider Q as a node depending on the node P . For example, when we say: *multiplication* depends on *addition*, it implies that we must learn/teach addition before multiplication. This is true whether we construe multiplication and addition as either concepts or operational skills (activities). *Division* depends on *multiplication*. Transitively, *division* depends on *addition*; *fraction* depends on *division* as well as *multiplication* and so on. By applying transitivity, we can infer that fraction depends on addition, and therefore if we introduce fractions soon after teaching addition that will be a disaster for obvious reasons. Bridging and walking through all intermediate steps, and not skipping, will give us the best teaching-learning sequence. If we use computer science jargon, we need to find a longest path algorithm, instead of a shortest path algorithm to solve the problem.

Often while teachers and textbooks talk of prerequisites their focus is mainly on concepts, and skills are not carefully thought about, possibly because it is more challenging to make explicit what procedural skills and experience the learner must have before introducing new constructions. Mathematics teachers are on better ground here, since most of the time the emphasis is mainly on skills. In general, a lot of knowledge remains implicit in the activities we perform, and the meaning of the terms we use is often entirely grounded in metaphors and the metaphors in turn, in the activities our body performs.[12]

Currently this information is not available for easy search even for concepts, and least available for activities/skills, for mostly it exists only inside the minds of the teachers or in the pages of the text books and in the curriculum designs. What we can do is to *harvest* them from these resources and publish at a commonly accessible place. Such a resource, to our astonishment, did not exist on the Internet or in any compilations published in the printed form. Several semantic networks exist, and due to a spurt in the activities in both research and applications in the semantic web, multiple ontologies have been published (e.g. [13, 4]). But nowhere could we see any interest in mapping dependencies. In the current age of collaborative culture and social networking, this seemed strange! Using semantic networks, web portals are actively collecting information about geneology, who is a friend of who, what links to what, who reads

what book, and buys what kinds of consumer items, and so on. But, no online resource for knowing what depends on what is available. The Wikipedia has resources on each topic, and each topic has very useful metadata including the categories that it belongs to. However, a tag to specify prerequisites is not supported.

However, standardized machine based e-learning specifications have a scope for specifying prerequisites, see e.g. SCORM specification³. Using semantics to guide expert tutoring systems is proposed, e.g. see [5], which also uses SCORM's sequencing model. However, the information contained within these e-learning modules remains isolated and no mechanisms have been suggested to merge such metadata to create a globally useful network. Most of this digital e-learning material is also held tightly due to proprietary interests. All these situations prompted us to embark on the current project.

The simplest thing we propose is to *gather* every assertion of prerequisites from all sources, subjects, and store them in a single large knowledge base. As we can expect we will get a massive semantic network holding activities and concepts linked by their dependencies. Of course holding this network on a paper format, or as a book will neither be possible nor convenient—a possible reason why nobody ever built them already. We could however easily store such information in a computer as a knowledge base, collect them over the Internet, and retrieve the resulting dependency network from the servers as graphs for any node of our interest. Considering that each node is also a *learning objective* (LO), the recursively drawn graph will give us a *road map* of each LO. Convinced that such a resource is going to be very useful for education, we embarked on this project and launched as a globally accessible collaborative portal at <http://www.gnowledge.org/>. An example one such map is shown in Figure 1 displaying the prerequisites of the concept “fraction”. You can find several hundreds of such maps dynamically created from the live site as the community began to contribute the seed content. The maps that we see now are far from being comprehensive or correct, because more missing linkages are to be added, and incorrect links deleted.

³ Advanced Distributed Learning Initiative, a consortium of agencies interested in e-learning resources developed a standard for representing Sharable Content Object Reference Model (SCORM), published at <http://www.adlnet.gov>.

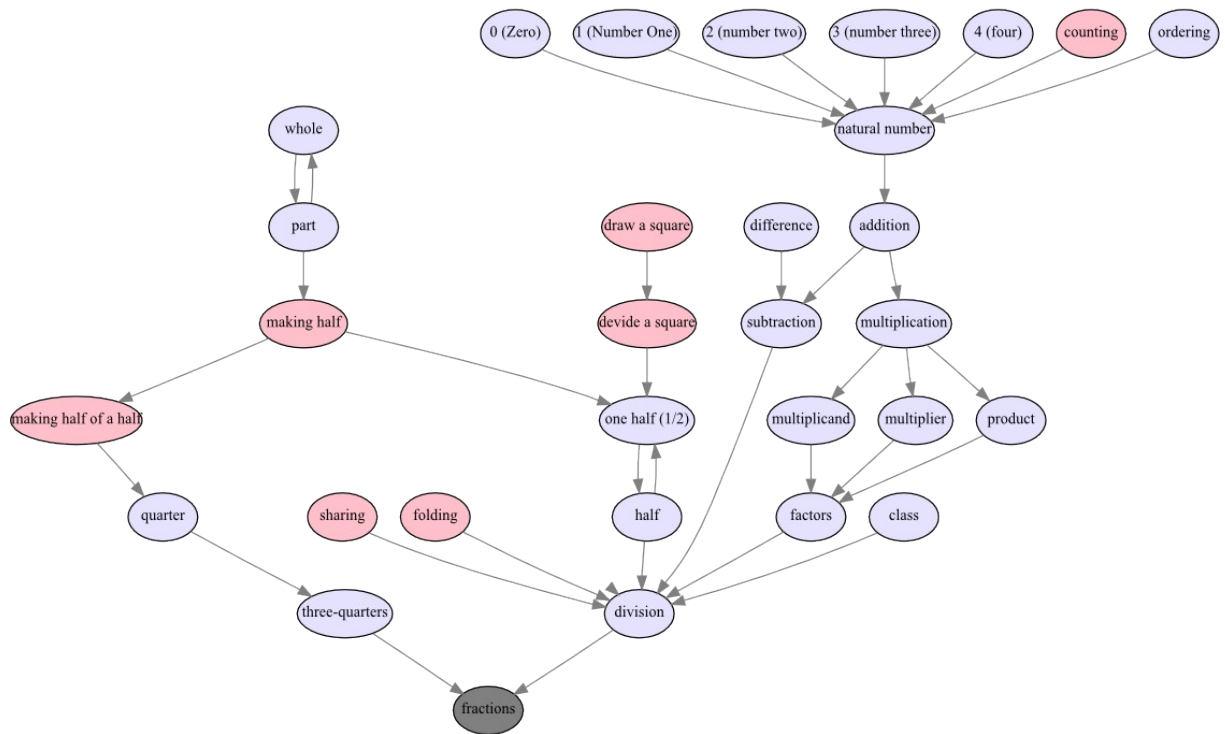


Figure 1: An example of a dynamically generated road map. The node for which the sequence is shown is coloured gray. The activities, in contrast to concepts, are represented in pink. Some of the nodes, e.g. part/whole, have mutual dependencies suggesting that their meanings are learned together in the same learning context.

The resulting network can generate different kinds of useful graphs:

1. the **road map** of any given learning objective depicting the multiple paths that *converge* to a destination as shown in Figure 1.
2. the **road-ahead map** of any given concept/activity, depicting the *divergent* routes one may explore after reaching a destination as shown in Figure 2.
3. the **dependency map**, which is a combination of 1 and 2 for a given node.
4. the **merged map**, a single directed graph of all the asserted links in the knowledge base as shown in Figure 4.

Though it may appear complicated, most of the complex job is done by the software, which is already in place, though we are constantly upgrading features and improving the algorithms for better performance as well as better visual appeal. Any learned community members, particularly educators, can submit the prerequisites as simple assertions as shown below:

[multiplication of fractions by whole number] < whole number; fraction; multiplication; denominator; numerator;

[division of fractions by whole number] < whole number; division; multiplication; denominator; numerator;

Avagadros principle < volume; gas; temperature; pressure; molecules;

average speed < speed; average; time interval;

Celsius scale < temperature scale; freezing point; boiling point; water;

sum < addition;

factors < multiplication;

subtraction < difference;

product < multiplication;

factors < product;

division < factors; product;

[checking division by multiplication] < division;

fraction < division; multiplication;

[solving division problems that result in remainder] < whole number; division;

division > remainder; divisor; quotient; dividend;

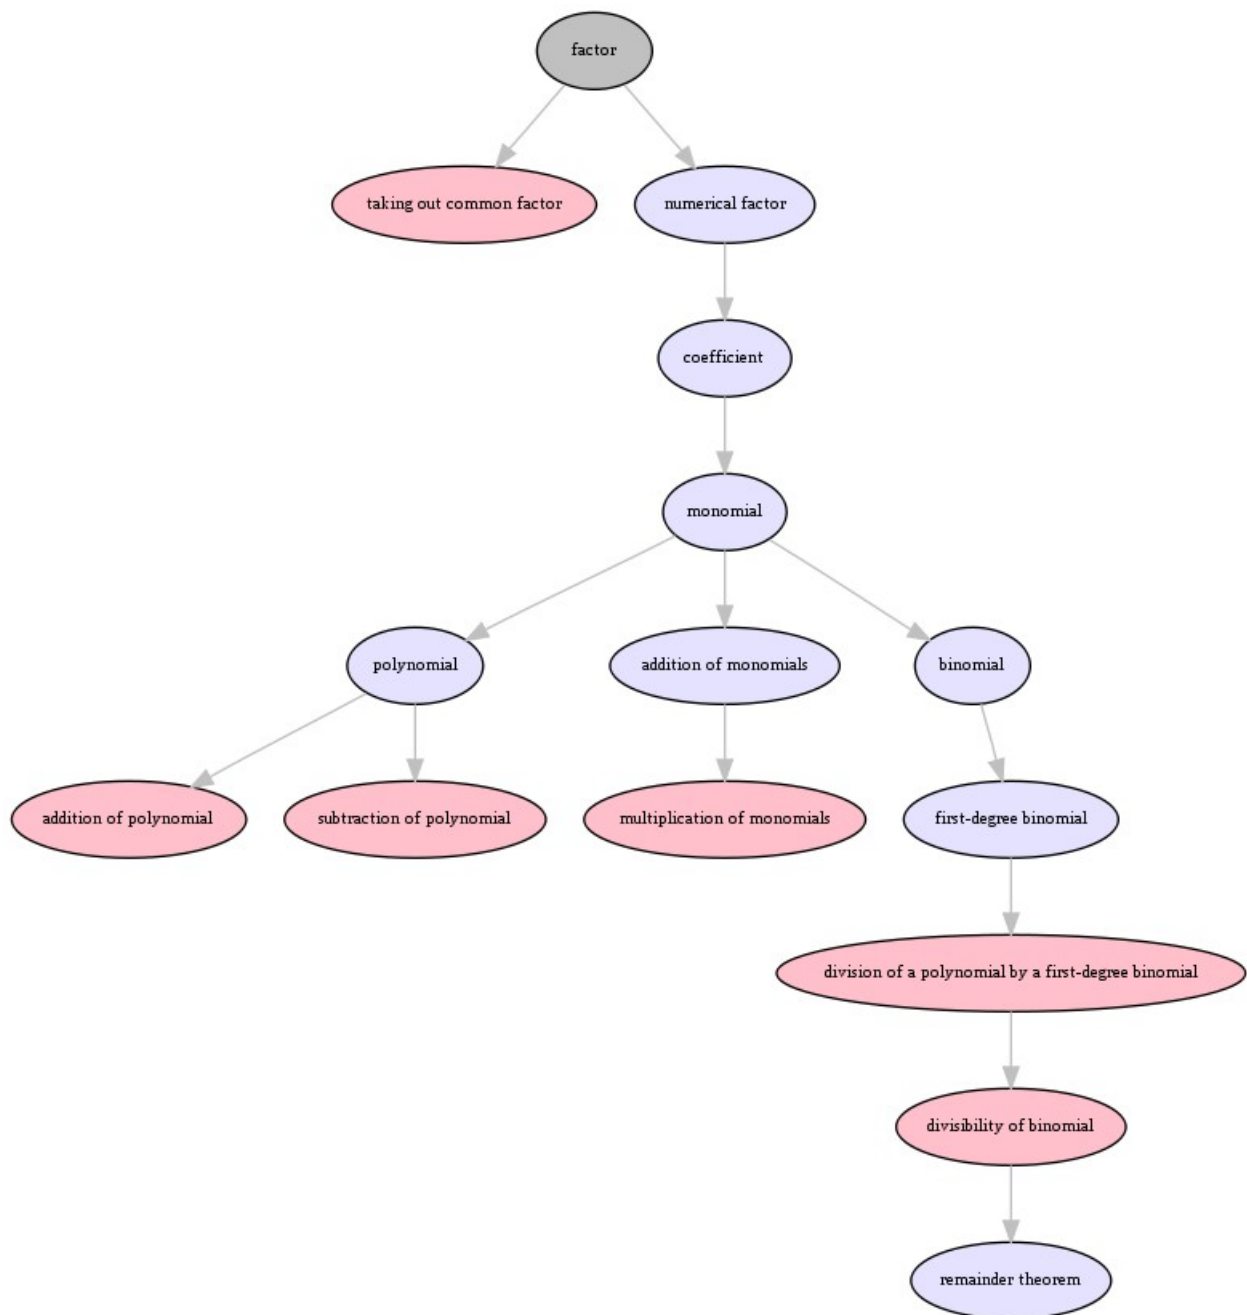
brackets > order of operations;

even number < divisability by 2;

These are examples of simple assertions of prerequisites that the users are expected to submit. Activities are distinguished by the use of square brackets “[]”. In the absence of an enclosing square bracket the expression will be inserted in the knowledge base as a member of a concept. The relation sign “<” is used to represent “depends on”, and “>” is used to represent its inverse “prerequisite for”. In a single sentence more than one dependency can be asserted by separating the expressions with “;”. The submissions can be in any order, and there is no need to segregate the assertions of one subject area from other.

Each user’s contribution is marked by their user name keeping a record of who contributed what and when. Every change is marked and recorded in the knowledge base, and nothing is permanently removed. There is also a provision to *delink* incorrect assertions, correcting spellings, or modifying the name of the expression etc. The authors can write to other authors to settle disputes in the mailing list. A feature to add a talk page, as in Wikipedia, is going to be made available soon, to make collaboration easy. A feature to merge nodes, translate to other languages, access to statistics, will be soon made available. The portal provides facilities to search and view the graph in several of the formats mentioned above. The site is constantly under development, with several new features getting added every month. Software developers can also contribute to the technical development of the portal. The sourcecode as well as the knowledge base are released under similar copyleft license, as in Wikipedia.⁴

⁴ A copyleft license is a creative and constructive use of copyright that grants the users certain rights (inplace of restricting terms imposed by most common use of the same copyright) such as: the permission to use for any purpose, make changes, share the changes with others provided the derived knowledge is also released under the similar copyright terms. See <http://www.gnu.org> for more details on how it works.



Map generated from knowledge.org on 16/7/2009 at 0:22 released under GNU FDL 1.2

Figure 2: The road-ahead map of the concept "factor" dynamically generated based on the data available as on 16/7/2009.

Having such a community portal will help not only the teachers, but also the researchers to critique and collaboratively update and arrive at consensus so as to arrive at a most authentic teaching-learning sequences. Just as several members of the community contributed in various ways and built a resource like

Wikipedia, we hope the maps at gnowledge.org could also be created, and will be equally relevant to the community.

In the next section the methodology of constructing the maps is elaborated along with a brief discussion of the possible didactic and epistemic implications.

2.2 The simple rules for making teaching-learning sequences.

The method of constructing the maps depicting the teaching-learning sequences is very simple, but humanly impossible without employing computers. The rules are as follows:

2.2.1 Each distinct node will occupy a unique position in the map.

The dependency network is a graph, therefore the collection all nodes, where each node is any unique concept or an activity, is a set and each node will have a single unique *relational* position in the map. Currently the rules apply for a two dimensional graph. The position of the node is determined by the longitude rule.

2.2.2 The rule for determining the position of nodes on the longitude of the map.

1. If Q depends on P , place P one level above Q .
2. If R depends on P , place R and Q one level below P .

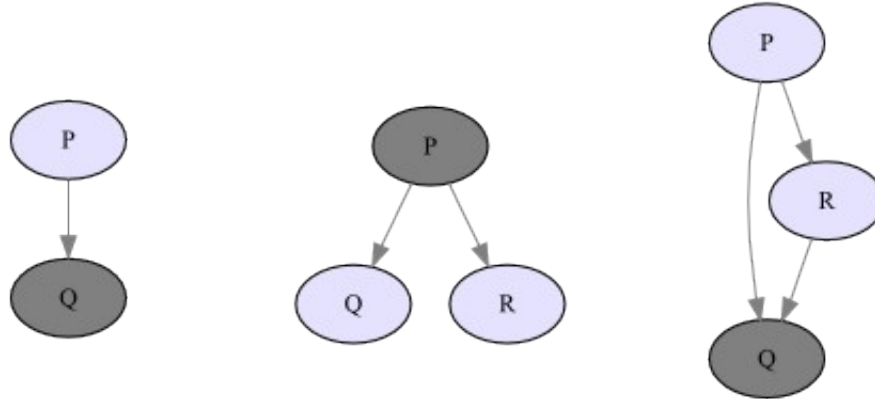


Figure 3: The Longitude Rule: Q depends on P , place Q below P ; Place Q and R in the same level, since both depend on P ; Push Q below R , since Q depends on R .

3. If Q also depends on R , then place Q one level below R , such that Q is two levels below P .

Apply this rule recursively till we obtain a graph where all nodes have unique position, and are located at different levels. Since this rule determines the vertical alignment of nodes, let us call this *the longitude rule*.

An immediate didactic implication of this rule.

All concepts and activities that are to be carried out at an early stage will automatically appear on the top layers of the graph. Subsequently, the bottom most layers are clearly the advanced concepts. This automatic layering can be put to immediate use for deciding the graduation levels while designing a curriculum.

An epistemic implication of this rule.

The longitude rule has the following important epistemic implication: the nodes are positioned according to how many degrees far are they in terms of their meaning providing a measure for *semantic distance*.⁵ In the above example, we can say that Q is separated from P by two degrees. This measure will depend and will change till we arrive at a comprehensive data set, and will eventually stabilize as a limiting case as the data increases. We leave this pointer without elaboration, since the topic of measuring semantic distance will

⁵ Measuring semantic distance is a classical problem in knowledge representation research and there are multiple approaches to solve the problem. Knowledge represented using dependency graphs as suggested here provides another approach. This proposal is a topic for another forthcoming article.

take us too far away from the current purpose.

2.2.3 The Rule for determining the position of nodes on the latitude of the map.

Consider a node P that depends on a number of nodes m , and is a prerequisite for a number of nodes n . By following the longitude rule, the m nodes will be above and n nodes will be below the node P . Following the vocabulary of directed graphs, we can call the links coming from m to P as *incoming links*, and the links going below to n from P as *outgoing links*, and the nodes m and n can be called *incoming nodes* and *outgoing nodes* respectively. Both outgoing nodes and incoming nodes can in turn have both incoming and outgoing links and so on. Let us define the total number of incoming and outgoing links of any given node as the *mass number* of the node. Now we can specify the latitude rule as follows:

The lateral position of the nodes is determined by a relative rule on the basis of the mass number of the nodes: the nodes with the lesser mass number will be pulled closer towards the nodes with the greater mass number. Though this is a simple rule, adding and removing links in the network will change the positions based on this rule. Nodes get pulled to the right or the left to each other as and when additions happen in the membership of the set. Since the form of the network changes whenever there is an addition or deletion of a link, we can also call this rule an *accommodation* or *alignment* rule. This rule has interesting didactic and epistemological implications.

Didactic Implications of the latitude or accommodation rule.

Nodes with large number of outgoing links cannot be neglected in education. At any given latitude, the first priority can be given to the nodes with larger number of outgoing links. Since, outgoing links are required for learning several other LOs, this priority makes sense. This provides the learner a *potential* to move on.

The LOs that appear in the same latitude must have similar degree of difficulty. When new links are added the membership of nodes in a latitude may change. Therefore, the greater the comprehensiveness of the graph, the greater will be the accuracy of the statement about the degree of difficulty. A few levels of proximate latitudes can be grouped together as a level of graduation, since they suggest a similar degree of difficulty.

Different disciplines (subject areas) get aligned automatically. Since we do not specify which concept belongs to which category or subject area, this result is very interesting and may have serious implications to knowledge representation research and epistemology in general and didactics in particular. The topics that are close together form, as it were, into islands, which can be distinguished as different subject areas based entirely on semantic dependency. Subject areas that have interdisciplinary relations are also seen close together. Since it is not possible to meaningfully present the already very large resulting graph, please see the current state of map, which looks like a milky way, published at <http://www.gnowledge.org/Data/map/mergedMap.png>. In the Figure 4 you can see the thin central spread of dots which contain about 2300 nodes.

Epistemological implications of the accommodation rule

Much of these epistemological implications I draw below are preliminary hypotheses and need greater elaboration and research. I state them anyway, for these foreseen implications might invite discussion, and strengthen the need for participating in the project.

Given a merged map of dependencies we can determine the nodes that have a greater number of incoming links or outgoing links than others. The nodes with highest number of either type of links will be very few in number, while the nodes with intermediate number of links will be quite high, and those that have very small number of links will be highest. This assortment is indeed a characteristic feature of all complex systems demonstrating the scale-free character.⁶ Several studies have already shown that the semantic networks of all natural languages do show the same scale-free character, and most of them are determined to follow a power law also known as Zipf's law [21]. This result indicates that:

1. The nodes that do not have incoming relations will be few and will be more than one. There will be no epistemic *summum bonum*, an ultimate concept from which all meaning will emanate, or an ultimate concept where all meaning will flow towards. Based on the logic one may predict that no node except the top nodes will have less than 2 incoming links. This is because logically fixing any meaning will need at least two concepts.⁷

⁶ An accessible and authoritative statement of this can be read from [3].

⁷ Philosophical literature is rich on discussions on this subject of individuating objects, events, actions and concepts. Strawson, Quine, Davidson etc. have elaborate discussions on this important subject. Since we have no space to dwell




Figure 4: Merged Map: The horizontal spread of dots in the middle is the shape of the graph containing about 2300 nodes with about 3600 relations. The data is not sufficient to get a true picture of the map as yet, but the map already shows related concepts together, and is already good enough to begin measuring semantic distance. The map shows some regions that contain neighborhood of the same subject. The zoomed-in graph on the top is the dependency graph of biological cell, while the bottom graph is that of water. All the graphs are generated after applying transitive reduction algorithm. High resolution maps can be viewed from the site. Panorama view of the merged map created on 13th July 2009.

2. The top nodes without incoming nodes, therefore, cannot be concepts, they must be nodes describing an *action*. This follows since while fixing any concept requires more than one term, an action on the other hand may be considered an innate ability of the learning agent, and its meaning can be considered implicit, and thus does not depend (explicitly) on any other. If true, this implication will also suggest that all meaning begins from the top activities. Prima facie support to this is already available from what we teach during the first few years after the child is born. The dependency map of that period will be full of activities, some of which eventually help in grasping the use of terms that stand for concepts. Therefore, we can predict that the top nodes of the map will be predominantly of activities.
 3. Though the current state of the merged map at gknowledge.org does show some disconnected
- on these reasons we let it pass.

islands without any connections to the rest of the network, if we look at the names of the nodes, it is very clear that they are disconnected only because no one has so far asserted the dependency relations, and not because no relations are semantically applicable. We hypothesize that *in a comprehensive dependency network no orphans are possible*. If this turns out to be true, the result can have serious implications in epistemology in general, and philosophy of science in particular. One of the major implication is that given any two concepts, it is possible to measure the semantic distance, irrespective of any area of knowledge. That is, *incommensurability* in a network of meaning makes no sense. Considering the major debates that took place in the history and philosophy of science following the famous Kuhn’s assertion of incommensurability between two different world views, two languages etc., the result that we are heading towards will have serious implications.[11, 16]

2.2.4 The transitive reduction rule generates the road map.

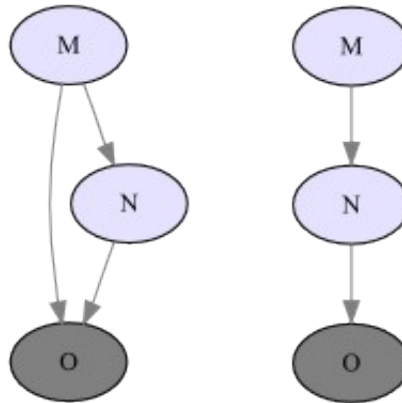


Figure 5: The graph on the right is a result of the transitive reduction rule.

Consider a concept M which is required for learning another concept N . And N is a prerequisite for another concept O . Though by applying transitivity one may say that M is required for O we cannot skip N in between. Thus given a dependency graph teaching-learning sequences cannot skip the nodes in between and go for a shortcut, as mentioned earlier. Therefore, the road-map for teaching-learning is not the shortest path in a dependency network, but the longest path. Using this logic we have applied a *transitive reduction*

algorithm, that removes the ‘skipping’ links while making the graph.[1] The user’s assertion is not deleted from the knowledge base, but while making the road-map the short-cut links are removed to get the complete teaching-learning sequence. This gives the possible path of semantic flow. See Figure 5 to find a comparison of graphs obtained with and without transitive reduction algorithm.

The obvious use of the road map obtained by transitive reduction for any LO can inform the learner/teacher what to do next. The teaching-learning sequence can also effectively be used by online e-learning environments to provide an automatic navigation mechanism. E.g, Wikiversity could ensure that the learner worked out all the prerequisite pages before attempting a lesson.⁸

3. The Making of an Atlas

Knowledge cartography is one of the areas of current interest by several research groups.[17] Though the project proposed in this paper is much in line with this trend, to the best of our knowledge there is no prior work on using the dependency links to create a single merged map of all knowledge. Since all learnable concepts and skills can find a place in the dependency map, where each node can have a determinate position with longitude and latitude of the map also specifiable, it is clear that a surface map of all knowledge is indeed possible.

Given such a surface map, various other semantic relations can be asserted on top of it, creating multiple layers one for each relation. For example, visitors could soon add relations expressing what *is a part of* what, what *causes* what, what *produces* what (as in products/reactants), what *is an assessment for* what, what *is prior-state* and *post-state* of a an event/process etc. Since the subjects of these other relations will also have to be the nodes in the dependency relation, all other assertions of knowledge must lie only on top of the dependency relation. Thus we foresee the possibility to make an entire atlas of all knowledge containing explicitly represented form of knowledge. As in a geographic atlas, each map shows the same surface in a different way (political, physical, geological, climate etc.), the knowledge atlas will provide multiple layers of presentation depending on the semantic dimension chosen.

Thanks to the increased use of free software and free knowledge initiatives around the globe, a large

⁸ Wikiversity <http://en.wikiversity.org>.

number of explicitly asserted semantic relations, other than dependency relation, are already harvested by projects like DBpedia [4], which by now have 4.7 billion assertions comprising multiple languages and ontologies, there is hope: an atlas of knowledge is not an impossibility!

In our earlier studies on *refined concept mapping* [8], [10] and *the roots of rigor* [9] the focus was on the minimal use of relation names for promoting understanding in science. Availability of most widely used relation names in a merged ontological space, as in the atlas or DBpedia enhances the possibility of further facilitating the refined concept mapping for science education. We expect such an encyclopedic resource will guide students and teachers both visually and operationally.

In another study we have extracted the dependency relations from the repositories of the largest operating system distribution (Debian GNU/Linux). The system has about 23,300 packages with about 99,000 dependency relations asserted among them. The results suggest that the semantic system based on dependency relation exhibits the properties of a complex system [15]. We have also embarked on the project, as an extension of the study, to extract every dependency of all functions defined in the entire GNU/Linux operating system. An artificial operating system will not work unless all meaning (operations) is explicitly stated. This artificial semantic system can act as a control for the natural human semantic system. The study of the semantic flow of this artificial system gave the encouraging results that added up to the confidence that similar results will be obtained in the case of human semantic system as well. Much of human knowledge is tacit [18], and much of the meaning is grounded in subjective experience making it very difficult to declaratively express every bit. However, the current project of making dependencies explicit is not required to be about the exact nature of meaning of every term or activity, but only an exercise to do dependency sorting. This seems to us an eminently feasible exercise.

4. Conclusion and Invitation to Contribute

In this short communication, a simple method of creating teaching-learning sequences is presented, and how it may lead to the creation of an atlas of knowledge. The task may take time, but the methodology is simple and seemingly effective. There will be technical problems along the way while handling millions of relations in a single repository, scalability and performance issues, and most importantly presentation of the relations without a clutter etc. Given the rising collaborative culture and increasingly expanding

cyberspace, we hope there will be a stigmergic effect: each of us while entering the simple assertions into the knowledge base may not even foresee what global effects we may likely see. This therefore is a worthy experiment to conduct. Please visit <http://www.gnowledge.org/>, where the atlas is being built, to participate in any which way: contribute relations, translate to other languages, point out errors, add assessment nodes for each resource, delink incorrect assertions, discuss the results, contribute software and algorithms, and most importantly test the teaching-learning sequences in actual practice to verify if this is indeed a worthy thing to pursue further.

Acknowledgements

The inspiration to use dependency relation and the faith in collaborative spirit rose from the manner in which the largest software distribution, Debian GNU/Linux operating system, manages installation of software packages by keeping a database of assertions specifying which package depends on what. The current project ideas were a result of a desire to build a Debian like free knowledge distribution for promoting free education. The project is supported by the Homi Bhabha Centre for Science Education, TIFR, India, under the XI Plan project. The team members of the gnowledge.org laboratory, Alpesh Gajbe, Arnab K. Ray, Divya, Ganesh Gajre, Jay Mehta, Meena Kharatmal, Rajiv Nair, Sashwat Chakraworthy, among many other past members of the lab, worked hard on the project.

References

[1] A. V. Aho, M. R. Garey, and J. D. Ullman. The transitive reduction of a directed graph. *SIAM Journal on Computing*, 1(2):131–137, 1972.

[2] David Ausubel. *The Psychology of Meaningful Verbal Learning*. Grune & Stratton, Oxford, England, 1963.

[3] Albert-Laszlo Barabasi. *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*. Plume Books, 2003.

[4] Christian Bizer, Jens Lehmann, Georgi Kobilarov, Soren Auer, Christian Becker, Richard Cyganiak, and Sebastian Hellmann. Dbpedia - a crystallization point for the web of data. *Journal of Web*

Semantics (to appear in a forthcoming issue).

- [5] Yu-Liang Chi. Ontology-based curriculum content sequencing system with semantic rules. *Expert Systems with Applications*, 36:7838–7847, 2009.
- [6] Rosalind Driver, Hilary Asoko, John Leach, Eduardo Mortimer, and Philip Scott. Constructing scientific knowledge in the classroom. *Educational Researcher*, 23(7):5–12, 1994.
- [7] Peter Hylton. *Quine*. Arguments of the Philosophers. Routledge, New York, 2007.
- [8] Meena Kharatmal and Nagarjuna G. A proposal to refine concept mapping for effective science learning. In A. J. Canas and J. D. Novak, editors, *Concept Maps: Theory, Methodology, Technology: Proceedings of the Second International Conference on Concept Mapping*, pages 1–7, San Jose, Costa Rica, 2006. Universidad de Costa Rica.
- [9] Meena Kharatmal and Nagarjuna G. Exploring roots of rigor: A proposal of a methodology for analyzing the conceptual change from a novice to an expert. In A. J. Canas, P. Reiska, M. Ahlberg, and J. D. Novak, editors, *Concept Mapping: Connecting Educators. 3rd International Conference on Concept Mapping*, pages 391–398, Tallinn, Estonia & Helsinki, Finland, 2008. IHMC, Tallinn University, University of Helsinki.
- [10] Meena Kharatmal and Nagarjuna G. Refined concept maps for science education: A feasibility study. In K. Subramaniam and Anwesh Mazumdar, editors, *Episteme 3: An International Conference to Review on Science, Technology and Mathematics Education*, pages 76–80, Mumbai, India, 2009. Macmillan Publishers India Ltd.
- [11] T. S. Kuhn. *The Structure of Scientific Revolutions*. Chicago University Press, Chicago, 1970.
- [12] George Lakoff and Mark Johnson. *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. Basic Books, New York, 1999.
- [13] Douglas B. Lenat and R. V. Guha. *Buuilding Large Knowledge Based Systems*. Addison Wesley,

Reading, Massachusetts, 1990.

- [14] Joel J. Mintzes, James H. Wandersee, and Joseph D. Novak. *Teaching Science for Understanding: A Human Constructivist View*. Academic Press, San Diego, 2004.
- [15] Arnab K. Ray, Nagarjuna G, Rajiv Nair. Semantic structure and finite-size saturation in scale-free dependency networks of free software. Preprint <http://arxiv.org/abs/0901.4904>
- [16] Eric Oberheim and Paul Hoyningen-Huene. The incommensurability of scientific theories. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2009 edition, 2009.
- [17] A. Okada, S. B. Shum, and T. Sherborne. *Knowledge Cartography*. Springer, London, 2008.
- [18] Michael Polanyi. *Personal Knowledge: Towards a Post-Critical Philosophy*. The University of Chicago Press, Chicago, USA, 1958.
- [19] John F. Sowa. Semantic networks. In Stuart C. Shapiro, editor, *Encyclopedia of Artificial Intelligence*. John Wiley & Sons, Inc., New York, 2 edition, 1992.
- [20] Stella Vosniadou and William F. Brewer. Theories of knowledge restructuring in development. *Review of Educational Research*, 57(1):51–67, 1997.
- [21] G. K. Zipf. *Human Behavior and the Principle of Least Effort*. Addison-Wesley, Cambridge, Massachusetts, 1949.