ELSEVIER

Journal of Alloys and Compounds 279 (1998) 47–59

Journal of ALLOYS AND COMPOUNDS

Predicting ultra-hard binary compounds via cascaded auto- and hetero-associative neural networks

Stephen L. Thaler

Imagination Engines, Inc. 12906 Autumn View Drive, St. Louis, MO 63146-4336, USA

Abstract

When an auto-associative neural network is trained within any conceptual space, the many rules and schema embodied within that knowledge domain are encoded through its many connection weights and biases. For instance, if such a network's input/output exemplars consist of numerous formulas representing known chemical compounds, subsequent network training will produce connection traces that embody the many implicit rules governing the constraints between constituent elements and their allowed proportionalities. That is to say, the net has gained a statistical 'insight' into the patterns of bonding, valence, and charge balance that must be observed in theorizing new chemical compounds. If that network is now made chaotic by random perturbation of its processing elements and connection weights, the resulting network activations will represent the formulas of a wide variety of plausible compounds, many of which may be considered novel from the standpoint of network training. We therefore attain an all-neural search engine for generating a stream of plausible chemical possibilities. Adding subsequent 'policing' networks to associate these emerging chemical formulas with various physical and chemical properties, we are able to either filter for sought characteristics or alternatively, assemble expanding materials tabulations of potentially new compounds and their estimated properties. Here, we describe the theory, construction, function, and results of just such an autonomous materials discovery machine, tailored specifically to the search for new ultrahard binary compounds.

Keywords: Creativity machine; Autonomous discovery; Ultrahard materials; Neural networks; Neural network cascades; Auto-associative networks; Virtual input effect; Attractor basins; Neural network skeletonization; Schema elucidation

1. Introduction

If we bombard the input units of an associative network with random values, the network's output patterns tend to cluster around those of its known training exemplars. Such a process is commonly known as vector completion. If, on the other hand, the inputs of the net are zeroed, while random perturbations are allowed to 'hop' from one connection weight to another, we see a related phenomenon: Now network output activations visit a mixture of both known training outputs and slightly degraded forms that retain many of the implicit constraints contained within the original training set. The overall network process of responding to internal perturbations, as though some environmental feature is present at the net's inputs, is known as virtual input effect [9]. The driver of this process is coined 'internal vector completion' to differentiate it from the traditional completion paradigm of supplying incomplete or 'noisy' vectorial inputs to a net. Using either of these effects separately, or in combination, we devise an all-neural search engine that may freely roam through any conceptual space in a process tantamount to human internal imagery and brainstorming [8,11–14]. The utility of such a search engine is the rapidity by which the underlying associative net may be rapidly trained using only historical data. We thereby bypass the often tedious and time-consuming chores of assembling the underlying rule or model bases.

Supplying any number of critic networks that may supervise the stream of outputs produced by such a continuously perturbed net, we form an automated, all-neural search engine [15,16]. Recently such virtual machines have been coined 'Creativity Machines' and have been shown to perform optimally at mathematically reproducible regimes of internal perturbation. To illustrate the utility of such a Creativity Machine within a chemical knowledge domain, we consider an auto-associative network that has been trained through some standard paradigm such as backpropagation [5] using a training set of chemical compounds of the general binary form A_xB_y. Here, A and B represent some choice of elements, that ultimately constrains the allowed choices of stoichiometries x and y. The representation of compounds may consist of the ground state electronic shell configuration (i.e., its

0925-8388/98/$19.00
PII: S0925-8388(98)00611-2

core, s, p, d and f electron occupation numbers) of the constituent elements A and B, along with analog values for the stoichiometric subscripts *x* and *y*.

We note and emphasize that the trained net is not to be regarded as a simple mapping in which network inputs are simply copied at the net's outputs. Instead, the reader should focus on the resulting vectorial chemical representation encoded within the net's hidden layers and then decoded at the net's outputs. Referencing Fig. 1(a), we see that the constraint relations between all attributes (i.e., chemical elements and their stoichiometries) are now encrypted within the network's connection weights. Therefore, output nodes (shown at the bottom of the figure) may not arbitrarily activate. They must obey the constraints gleaned within the training and inherent to the implicit rules of chemical bonding between elements. Therefore, applying the electronic representations for hydrogen and oxygen to the inputs of the net, the most likely stoichiometric outputs of the net would be either $x=2$ and $y=1$ (i.e., H_2O) or $x=2$ and $y=2$ (i.e., H_2O_2).

Hence the utility of such an auto-associative net is its inherent ability to perform vector completion wherein incomplete or missing information at the network's inputs is filled in at its outputs. Therefore, as depicted in Fig. 1(b), if the network has been adequately trained, it may be presented with both the correct electronic representations for hydrogen and oxygen and chemically implausible subscripts of 1.3 and 0.8. Through its cumulative training, perhaps upon thousands of other binary chemical compounds, the network will supply realistic stoichiometries at its outputs, approaching values of two and one. Repeating this procedure, we could perhaps maintain hydrogen and oxygen as the input elements, but now apply different random stoichiometric 'seeds' of 1.8 and 2.5 for the *x*, *y* subscripts. Now the stoichiometric outputs may settle into states approaching two and two (i.e., H_2O_2). We therefore define a technique that allows us to predict multiple stoichiometries for the same input combination of elements A and B. The chemical formulas emerging from such a network obey the cumulatively gleaned statistical constraints of chemical stoichiometry, suggesting a wide range of neutral chemical species. Therefore we may see network outputs representing H_2O_2, but not OH (i.e., OH^-). Similarly, other perturbations may produce outputs representing Si_3N_4, that may have been a training exemplar, or C_3N_4 that may now be a generalization of the original training exemplars (i.e., it is isoelectronic with Si_3N_4).

In essence, each of the chemical training exemplars now forms a 'memory' within the network, represented by a deep attractor basin. By converting the autoassociative network of Fig. 1(a) into a recurrent net and then iteratively recirculating outputs back into inputs, the network will most likely fall into one of these attractor basins when presented with a random input seed (i.e., noise). Alternatively, the network could fall into a shallower attractor basin that is not representative of any of its training exemplars (i.e., a confabulation or false memory that embodies most of the statistical constraints of chemical bonding). By its very nature, the emerging chemical compound, say C_3N_4, is plausible, since the myriad constraints of the conceptual space are being exercised. Here, the net is 'seeing' the similar valence shell configurations for both silicon and carbon, ignoring the electron cores, to produce the most likely chemical subscripts for the compounds three and four.

As implied above, the real utility of the full auto-associative net is the ability to predict multiple stoichiometries for any given elements A and B. Therefore, pinning the inputs of the net to representations of hydrogen and oxygen, we may introduce random seeds for the stoichiometries *x* and *y*, Fig. 2(a) and b. Depending upon the values of the seeds, network activation will tend to bifurcate into the output states approaching those representing H_2O or H_2O_2. Successive recirculations of outputs back through the inputs will successively move the *x*, *y* subscripts quantitatively toward the attractors (2,1) or (2,2) hence generating two plausible chemical compounds. Repeating this procedure of supplying various elemental combinations A and B, along with random stoichiometric input seeds, and allowing the network to converge to an agreement between input and output vectors, we may produce myriad plausible chemical compounds. This device then constitutes a chemical search engine that readily 'imagines' candidate compounds whose physical properties

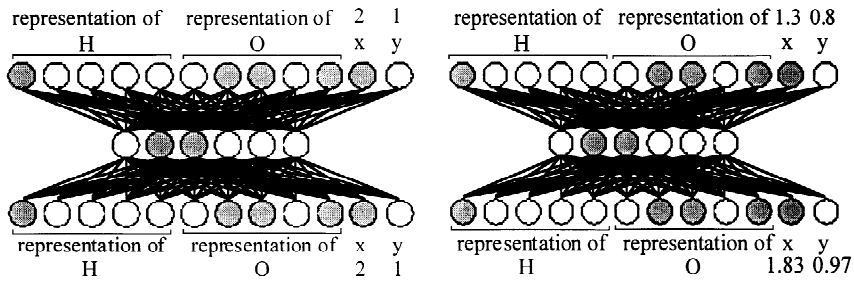

Fig. 1. (a) Auto-associative net mapping binary chemical compounds identically to themselves. Here, the application of the training exemplar H_2O to the net's inputs (top) produces H_2O at the net's outputs (bottom); (b) Application of incorrect stoichiometries *x* and *y*, along with the correct representations for the elements hydrogen and oxygen, produces corrected subscripts, representing H_2O. This filling in of incomplete or 'noisy' inputs is the process known as vector completion.

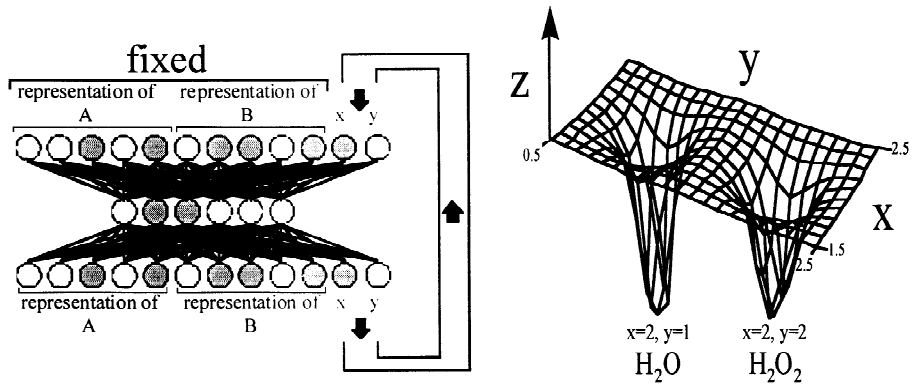

Fig. 2. (a) Converting the auto-associative net of Fig. 1 to a recurrent net, we may iteratively arrive at plausible stoichiometries *x* and *y* for any elements A and B. Convergence of the network to a potential chemical species occurs when the vector distance, *z*, between vector input (top) and output (bottom) falls below some predetermined threshold; (b) For the specific example of hydrogen and oxygen, the properly trained auto-associative network will possess two principal attractor basins corresponding to H_2O and as intimated in the idealized plot above. There, the *z*-axis corresponds to the vector distance between network inputs and outputs.

may be estimated by downstream hetero-associative networks. Here on, we use the term 'Creativity Machine' to represent this cumulative neural network cascade that first imagines chemical possibilities via such stochastic processes within its initial network stages, then passing these chemical candidates to downstream hetero-associative nets that now estimate a wide range of properties based upon these chemical formulas.

In a preliminary exercise of this materials Creativity Machine, a dynamic database has been generated consisting of potential binary compounds having the formula A_xB_y. The materials properties estimated by the downstream network modules within this machine have included a number of important physical characteristics including projected Mohs scale hardness. Following prolonged runs of this Creativity Machine, the resulting materials database, embedded within Microsoft Excel spreadsheets, could be ranked in terms of certain physical or chemical attributes. Concentrating on the predicted hardness of these materials, potentially new ultrahard compounds could be isolated, thus paving the way to a powerful materials discovery technique. Overall reliability of this materials search paradigm technique was evaluated in terms of (1) the plausibility of the emerging chemical compounds, (2) the accuracy of the resulting physical properties and (3) the relative proportion of new compounds generate above and beyond the initially available training exemplars.

2. The traditional path to ultrahard materials discovery

For comparison, perhaps the most recent and sophisticated example of the theoretical anticipation of new ultrahard materials stems from the work of Liu and Cohen [1,2] wherein a combination of empirical and ab initio approaches have been pursued. The overall pattern of discovery consists of an intuitive choice for the new ultrahard material based upon structural and electronic similarities with other known low compressibility solids. In the case of these researchers, the investigation of C_3N_4 was motivated by its close correspondence with the beta phase of Si_3N_4. Using a well-known empirical relation derived by Van Vechten [3] relating bulk modulus (i.e., its microscopic hardness) to bond length for tetrahedral solids, the corresponding beta phase of carbon nitride was expected to possess a bulk modulus approaching, or perhaps exceeding that of diamond.

Specifically, the Van Vechten equation takes the form

$$B = (19.71 - 2.20\lambda)/d^{3.5}, \quad (1)$$

where B is the bulk modulus in megabars, d is the bond length in angstroms, and lambda is a measure of ionicity of the compound. The constant 19.71 is inferred from the band gap calculated from reflection spectra, thereby qualitatively yielding some estimate of the energetic difference between the bonding and anti-bonding states of the material (i.e., the energetic difference between constituent atoms of both elements (1) chemically bound within a crystal lattice and (2) at infinite separation). Based upon the initial assumption that the carbon nitride structure would be tetrahedral and of low ionicity, initial estimates for the bulk modulus of this theoretical material exceeded that of diamond, thus motivating a series of highly ambitious band calculations requiring substantial computational resources. The ensuing computation showed a slight deviation from the initial structural assumption of a tetrahedral network to distorted derivative thereof, as well as larger than anticipated values of the ionicity parameter lambda. Still, based upon lingering uncertainties in the estimation of the ionicity parameter, there remained sufficient enthusiasm to motivate the synthesis and characterization of this material. These efforts still persist within various prestigious laboratories around the world. However, the lack of reproducible results in synthesizing this material suggests a possible

failure in anticipating the overall thermodynamic stability of the material.

In short, the search carried out by Liu and Cohen was tantamount to the proposed Creativity Machine Paradigm wherein a new compound and phase was generalized (i.e., imagined from past materials patterns by biological neural networks) from known materials and then evaluated for hardness using computer algorithms. This scheme is very similar to the intended connectionist plan to generate a host of materials, isoelectronic with known, characterized materials, and then to evaluate the hardness of each of these candidates compounds. A major difference is that the hardness evaluation will be carried out not by algorithmic means, but by a second neural network trained to recognize the underlying patterns of hardness as they relate to chemical formula. As qualification on the emerging materials, we will simultaneously predict thermodynamic stability (as well as other accompanying properties), through the free energy of formation, so as to eliminate the purely theoretical possibilities from those attainable within the laboratory.

3. Construction and operation of the materials discovery

3.1. Network training

Over a period of roughly 3 years, a database of over 20 000 inorganic compounds, each containing at most four distinct chemical elements, were collected from a variety of sources. The assembled information consisted of chemical formula, melting point, density, free energy of formation, and lattice constants gleaned from numerous X-ray crystallography databases. In building a separate hardness database, approximately 400 mineralogical examples were accumulated (largely from the 1996 CRC Handbook and mineralogical texts such as that by Pough, 1960 [4]). About 200 additional hardness values were assembled from the scattered literature on ultrahard materials, usually reported as values in Gpa or Knoop values. Since the majority of the data was mineralogical and reported in terms of Mohs scale hardness, small feedforward nets were trained to convert the Gpa and Knoop values to approximate Mohs values.

Once assembled into spreadsheet format, both databases were processed using a combination of parsing and conversion schemes that translated chemical formulas, represented in ASCII code, to numerical values representing the respective ground state electronic configurations of the respective component elements. Each element was represented by a vector consisting of five components, the number of core electrons (i.e., its inert gas core), and four numbers denoting the element's s, p, d and f shell valence shell populations. Therefore an element such as carbon was represented as the vector (2, 2, 2, 0, 0), with the first component representing the element's helium core ($Z=2$) along with 2 s and 2 p valence shell electrons. Each compound was represented by a supervector containing four subvectors representing the four possible elements, followed by four analog values representing the respective stoichiometric subscripts. Therefore, the formula for water, H_2O, would be represented through the vectorial notation ((0,1,0,0,0), (2,2,4,0,0), (0,0,0,0,0), (0,0,0,0,0), (2, 1, 0, 0)). The general template for representing a chemical compound took the form ((CoreA,As,Ap,Ad,Af), (CoreB,Cs,Cp,Cd,Cf), (CoreC,Cs,Cp,Cd,Cf), (CoreD,Ds,Dp,Dd,Df), (sA,sB,sC,sD)), with A,B,C and D denoting the four elements, and sA,sB,sC and sD, the respective subscripts of these elements within the compound formula.

In all, 11 separate feedforward nets were trained in constructing the planned cascade. The commercially available trainer Neuralyst™ was used in training these nets, using standard backpropagation, gradient descent methodology. Using Neuralyst's built in genetic supervisor, the hidden layer architecture and learning parameters were optimized for each network module using reserved test data. Below, we summarize the particulars for selected component network modules within the overall Creativity Machine cascade.

3.2. Element generators (network 1)

A three layer auto-associative net (Fig. 3), was trained to reproduce electronic configurations of the first 100 elements in the periodic table. Using the vectorial representation discussed above, training proceeded until outputs of the net agreed with its inputs to within 1% rms of the total range of each network output. This net was then converted to a spreadsheet network [12] by linking spreadsheet cells containing sigmoidal squashing functions by relative references weighted by the trained-in connection weight values. Linking outputs to inputs using a VBA

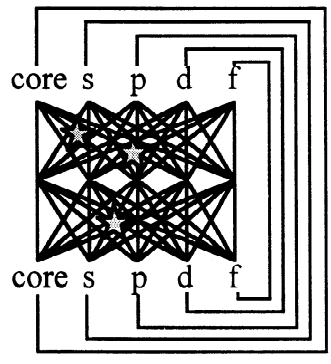

Fig. 3. 5/5/5 Recurrent network 1, used to generate elemental electron configurations. Any internal noise (shown as stars) introduced to the connection weights and successive passes through the net will result in vector completion and the generation of a valid elemental electron configuration.

(Visual Basic for Applications) patch, the feedforward net was converted to a recurrent architecture. The introduction of small perturbations to any of the network's connection weights would launch a succession of recursive cycles of that net, through which internal vector completion would produce a valid electronic configuration for some element, A. Fully exploiting the portability of the spreadsheet implementation, three copies of this recurrent net were cloned by copy and paste commands. All four copies of this net were placed side by side within the spreadsheet. Their role was to randomly generate elemental combinations A, B, C and D, thereby driving the Creativity Machine search and initiating the propagation of stochastically generated chemical information through the cascade.

3.3. Stoichiometry generator (network 2)

A three layered auto-associative net was likewise trained to within 1% rms error of each output parameter's range, Fig. 4. Input and output vectors consisted of only valence shell configurations and not the full electron configurations output by network A. This architectural decision was taken to calculate stoichiometry purely by generalization to similar, isoelectronic configurations within the training set. Therefore, in this scheme, C_3N_4 and Si_3N_4 would appear indistinguishable at this level of cascade processing. Once trained on over 10 000 chemical compounds, this net was likewise converted to the spreadsheet format. Appended to each of the connection weights within the Excel formulation, was a small chaotic term whose amplitude depended upon the estimated charge of the imagined chemical species (calculated by network 3, below). Therefore, as long as the net calculated charge on the candidate compound was non-zero, recursion between outputs and inputs could not settle within an attractor basin. When a neutral chemical species was finally imagined by network 2, the compound representation was made available for further propagation to downstream hetero-associative nets that evaluated various physical properties (nets 4–11).

It is to be noted that nodal architecture chosen for this network was determined by minimization of training error (in contrast to testing error). This unusual deviation from standard operating procedure was followed in order to assure that the auto-associative net memorized all 10 000 chemical training patterns. Such an approach is consistent with Creativity Machine design as prescribed in Thaler, 1997 [15] and with the underlying Creativity Machine theory [16], wherein useful ideas (i.e., novel chemical compounds) are nothing more than degraded memories induced by internal damage or noise within the network.

3.4. Charge evaluator (network 3)

A three layer feedforward network (Fig. 5) was trained to 3% rms of its total output range of nine charge units. The training set consisted of approximately 3000 exemplars, approximately 800 of which were known ionic species. The remainder of the training set was made up of known neutrals borrowed from the 10 000 compounds used for stoichiometry training. Inputs to the net consisted of the valence shell configurations of each constituent element along with the *x* and *y* stoichiometries. In Fig. 6, a correlation plot for the trained network contrasting network charge prediction with the actual charge is shown.

3.5. Free energy evaluator (network 4)

A three layer network (Fig. 7) was trained to within 6% rms error for the total free energy range of the training exemplars, -450 to $+50$ kcal mole^{-1}. Inputs consisted of the electronic representation of the four element compound and accompanying stoichiometries *x* and *y*. A correlation plot showing the accuracy of network predictions of free energy of formation is shown in Fig. 8.

3.6. Mohs scale hardness evaluator (network 5)

A three layer network (Fig. 9) was trained to within 3% rms error in mapping binary compounds to their anticipated Mohs scale hardness. In all, 339 training exemplars were used in training this net, chosen largely from

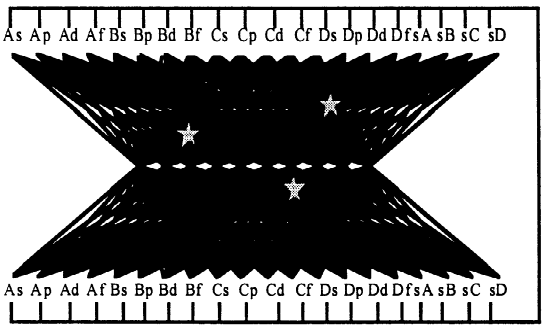

Fig. 4. 20/10/20 Recurrent network 2, used to generate plausible stoichiometries. s, p, d and f electron configurations representing as many as four separate elements, are generated in the preliminary network modules and transmitted to inputs of this net. Following successive recursions, outputs sA, sB, sC and sD yield plausible stoichiometry for the compound $A_{sA}B_{sB}C_{sC}D_{sD}$.

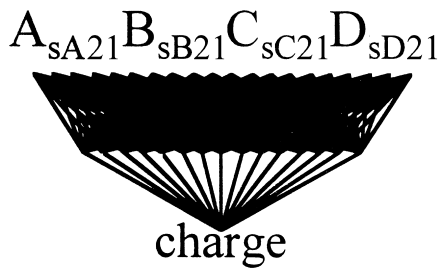

Fig. 5. 24/15/1 Feedforward net 3, used to evaluate charge on chemical species imagined by network 2.

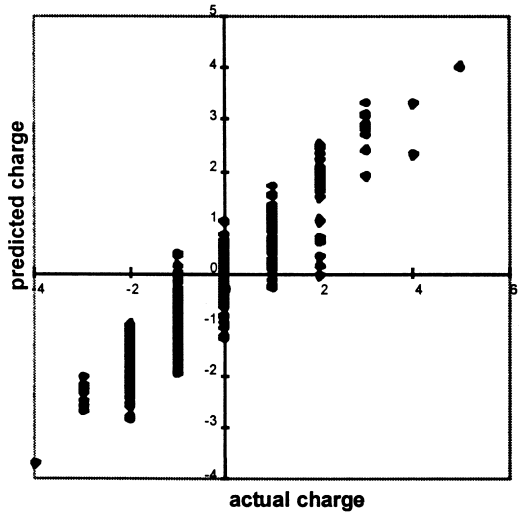

Fig. 6. Predicted versus actual charge for the fully trained network 3. Twenty percent of the data shown in the plot represents reserved test data.

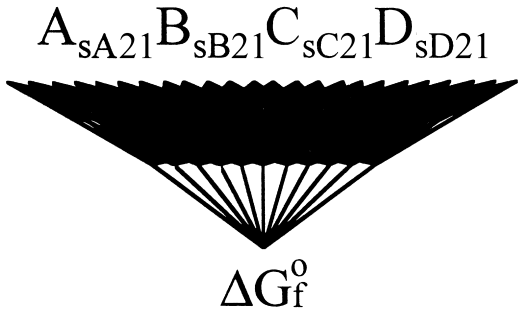

Fig. 7. 24/11/1 Feedforward net 4, used to evaluate free energy of formation of the imagined compound.

mineralogical references. If for any given compound within the net's training set there arose a choice of phase, the hardest of the room temperature and atmospheric pressure phases was chosen. Therefore, within this training

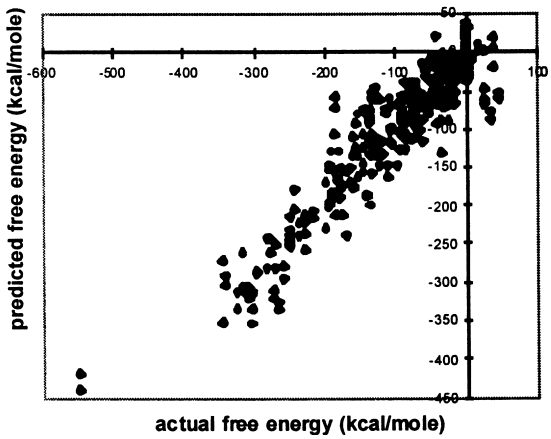

Fig. 8. Predicted versus actual free energy of formation for the fully trained network 4. Twenty percent of the data shown in the plot represents reserved test data.

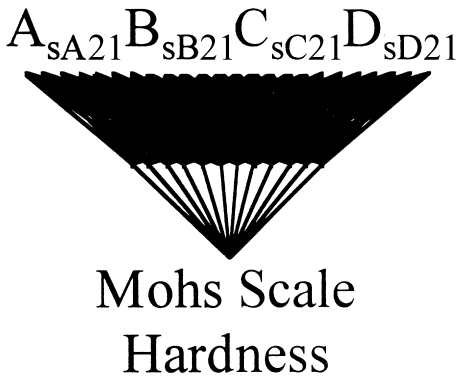

Fig. 9. 24/12/1 Feedforward net, used to evaluate free energy of formation of the imagined compound.

set, carbon is represented as diamond and not graphite, Fig. 10.

3.7. Remaining hetero-associative nets

In a similar fashion to the nets used to calculate hardness and free energy of formation, separate network modules were trained to relate chemical formulas emerging from network 2 to; (1) melting point, (2) density and (3) lattice parameters gleaned from X-ray crystallographic data. Training error for each of these nets was maintained below 10% rms error.

3.8. Cascade construction

All 11 of the network modules discussed above were converted into Excel spreadsheet format and pasted within an Excel Worksheet in the relative positions shown in Fig.

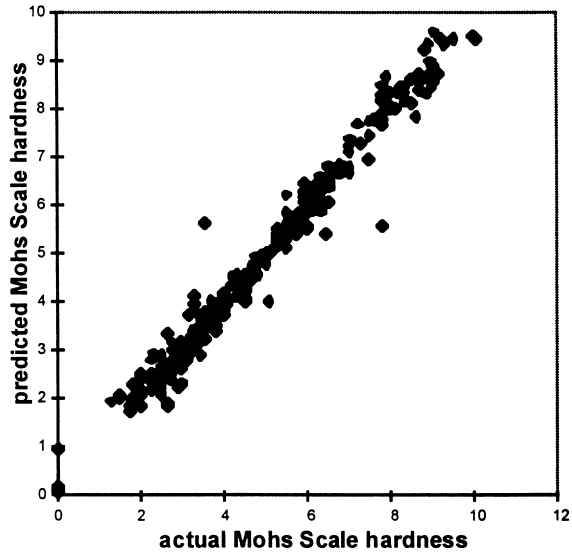

Fig. 10. Predicted versus actual free energy of formation for the fully trained network 5. Twenty percent of the data shown in the plot represents reserved test data.

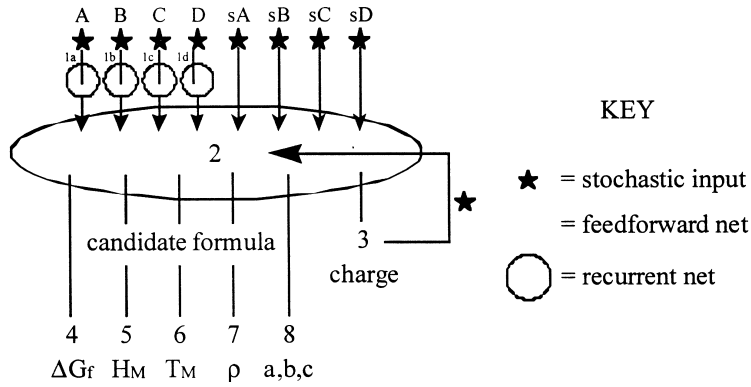

Fig. 11. Overall cascade structure showing placement of the trained network modules discussed above. The stars indicate stochastic inputs to the four clone networks generating plausible elemental electron configurations as well as random seeds for the initial guess of stoichiometry.

11. Where stochastic noise sources were required, as indicated by the stars, cells were filled with the resident Excel function rand(), yielding pseudorandom numbers between zero and one, the typical input range for computational neurons. Thus, noise terms were applied to the connection weights within the four networks labelled 1a–d, corresponding to the four initiating elements A, B, C and D. Likewise the inputs of network 3, determining net charge on the imagined species were connected to the outputs of network 2, containing the electronic configuration and stoichiometries within each imagined compound. The output of net 3, the estimated charge, was then multiplied by random factors supplied by the random function and then distributed by relative reference to each of the connection weights in network 2. All of the outputs of network 2 were manually connected by relative reference to the inputs of all the downstream network modules used to calculate materials properties.

3.9. Cascade operation

Applying perturbations to the connection weights of all four initiating networks labelled 1a–d, each activated into states representing the ground state electronic configurations of four randomly chosen chemical elements. In this study, however, we utilized only two of these networks, representing elements A and B, respectively. The outputs of the remaining two networks, normally corresponding to elements C and D, were pinned at values of zero. (Future runs of this Creativity Machine will allow similar stochastic choices of elements C and D to form potential quaternaries).

The first two stoichiometry nodes were set to random values, while the last two of this group were set to zero to indicate the absence of elements C and D. Network 2 was then run recurrently until the vector distance between normalized input and output vectors fell to 0.001. At each cycle of the recursion, network 3 estimated the charge on the chemical species appearing at the outputs of network 2. As long as this calculated charge approached a value of zero (i.e., a neutral species), the perturbations to the internal connection weights likewise ceased, allowing the now quiescent net to fall into one of its attractor basins, and hence into a state representing a valid chemical compound. If the estimated charge was non-zero, the introduced chaos within network 2 prevented it from recurrently converging to consistent output and input vectors.

Having generated a plausible chemical formula, the representation of that formula was then passed to each of five hetero-associative nets that estimated values for standard free energy of formation, Mohs scale hardness, melting point, density, and the lattice constants *a*, *b* and *c*. It was at this point that the generation cycle was reinitiated, beginning with the stochastic activation of the network 1 clones.

Following each cycle of compound generation and properties evaluation, the imagined compound A_xB_y was then submitted and compared with a static database of 4000 known compounds that had been reserved for network testing. If the imagined compound matched a formula within this reserved database, the network determined subscripts *x* and *y* were logged along with the actual subscripts into a table that was then later used to examine the correlation between predicted and actual stoichiometries. Also, as part of the runtime diagnostics, we have tracked the total number of compounds imagined, the number of compounds falling outside of the training set (i.e., invented, in contrast to memorized compounds), as well as a number of cross checks to examine the validity of the cascade properties predictions.

4. Results

In this initial run of the materials Creativity Machine, roughly 81% of the imagined compounds were found to be distinct from those contained within the initial training set of 10 000 compounds. While we cannot be perfectly assured of the plausibility of all compounds imagined, we

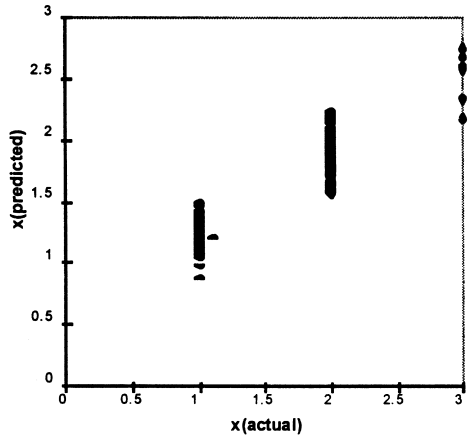

Fig. 12. Predicted versus actual *x* stoichiometry for 1000 of the predicted compounds.

may contrast the stoichiometries predicted by the imagination engine and contrast them against corresponding instances found within a reserved database of inorganic chemical compounds. Figs. 12 and 13 show that the

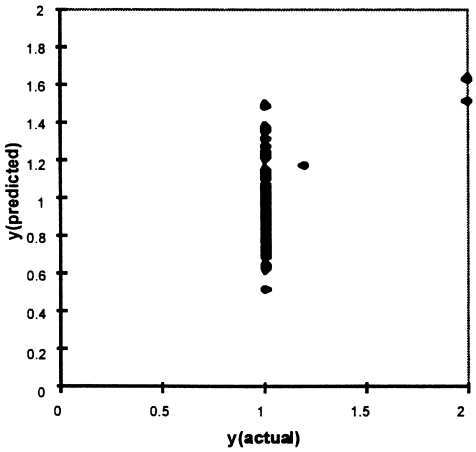

Fig. 13. Predicted *y* versus actual *y* stoichiometry for 1000 of the predicted compounds.

imagination engine is realistically predicting the stoichiometry of 'unseen' compounds. We note of course that there is an inherent spread in the recommended subscripts with slightly less than quantitative predictions. This qualitative prediction of chemical formula may therefore be thought of as a form of fuzzy chemistry that would help the experimentalist to zero in on promising chemical possibilities.

We show below the resulting data for the hardest of the materials generated during the autonomous run of the Creativity Machine. These consist primarily of carbides, borides, and beryllides, silicides, and surprisingly, a host of intermetallics, as shown in Tables 1–5, where we show some of the hardest materials within each grouping. In each of these tables, the candidate compounds have been ranked in ascending order by free energy of formation. Therefore, materials at the top of each table are deemed more plausible in terms of chemical stability.

In interpreting the stoichiometry of each of the predicted compounds, we note that when the generated compound matches some entry within the reserved materials database, subscripts have been converted to the integer values (i.e., $C_{1.00}$ $S_{1.00}$ in Table 1) encountered there. However, in the case that the imagined compound could not be found within the reserved data base, the fuzzy chemical subscripts have been retained. Therefore, a compound of the form $A_{0.75}B_{1.00}$ could be interpreted as either a fractional stoichiometry material or, rounding subscripts to appropriate integer values, as A_3B_4.

The hardest and most stable phase projected by this Creativity Machine run may be seen in the carbide listing of Table 1, as the entry $C_{1.00}C_{1.00}$, symbolizing in our representation the pure element carbon. We note however, that this material may not represent the graphitic or diamond phase, owing to the large negative free energy of formation (i.e., these values are 0 and $+0.64$ kcal mole^{-1} for graphite and diamond, respectively), and the low space symmetry suggested by the projected lattice constants *a*, *b* and *c*. This material may instead represent a low symmetry

Table 1
Ultrahard carbides imagined by the Creativity Machine and ranked by descending chemical stability

A	*x*	B	*y*	ΔG_f^0 (kcal mole^{-1})	MP (°C)	Mohs Hardness	Density (g cm^{-3})	xtala(Å)	xtalb(Å)	xtalc(Å)
C	1.72	H	1.01	−122.14	2172	9.71	2.29	9.01	8.55	6.15
C	1.00	C	1.00	−120.55	2922	9.80	2.36	4.98	5.20	3.94
C	1.00	Si	1.00	−95.74	2979	9.75	2.55	5.53	5.72	4.09
C	2.00	Mg	1.00	−93.46	1802	9.76	3.02	4.64	5.00	3.98
C	1.71	K	1.01	−86.7	1223	9.72	2.80	9.37	8.98	6.46
B	1.73	C	1.35	−75.9	2392	9.71	2.71	4.75	5.03	3.90
C	2.00	Ca	1.00	−75.87	1859	9.75	3.25	5.27	5.61	4.17
C	2.00	Ra	1.00	−45.42	1915	9.74	8.60	8.19	8.48	5.30
C	1.00	U	1.00	−26.04	2567	9.72	13.24	6.43	6.78	4.21
C	1.88	Np	1.13	−25.11	2517	9.81	12.39	6.86	7.13	4.53
C	1.88	Pu	1.13	−21.18	2653	9.82	13.06	6.66	6.93	4.41
C	1.87	Cm	1.11	−10.6	2902	9.77	16.58	5.52	5.87	3.77
C	1.87	Cf	1.11	−4.27	3070	9.78	17.61	5.46	5.76	3.65

Table 2
Ultrahard borides imagined by the Creativity Machine and ranked by descending chemical stability

A	x	B	y	ΔG_f^0 (kcal mole^{-1})	MP (°C)	Mohs Hardness	Density (g cm^{-3})	xtala(Å)	xtalb(Å)	xtalc(Å)
B	1.73	C	1.35	−75.90	2392	9.71	2.71	4.75	5.03	3.90
B	1.00	H	1.00	−50.35	1301	9.78	2.31	8.60	8.22	5.58
B	1.00	Li	1.00	−46.20	1355	9.77	2.33	8.69	8.32	5.66
B	2.00	Be	1.00	−38.69	1527	9.74	2.46	4.28	4.66	3.67
B	1.00	Na	1.00	−30.09	1619	9.69	2.45	8.94	8.60	5.89
B	1.93	Pb	1.74	−6.75	831	9.79	11.24	8.34	7.90	3.57
B	1.00	Rb	1.00	2.93	1483	9.69	3.35	8.50	8.39	6.10
B	1.94	Bi	1.86	4.25	1724	9.77	11.71	9.39	8.81	3.72
B	1.93	Np	1.18	19.25	2252	9.75	11.32	6.96	7.27	4.20
B	2.00	Pu	1.00	19.75	2262	9.77	12.00	6.75	7.04	4.12

phase of pure carbon, perhaps some distorted form of bridged graphite [7]? Another potential carbon-based ultrahard compound shown in Table 1, is $C_{1.72}H_{1.01}$ (perhaps C_2H_1), perhaps intimating in general that hydrogen-deficient carbon networks (i.e., a – C:H) may in fact approach diamond in hardness.

One of the more plausible, ultrahard borides encountered in Table 2, is $B_{1.00}H_{1.00}$. Such a material would be analogous to the hydrogen-deficient carbon networks discussed above. In contrast to the case with carbon, however, hydrogen would no longer be terminal, but a bridging atom between borons. The additional bond energy per unit volume would then significantly contribute to the overall hardness of this material.

We may also generalize this potential crystalline phase to a family of disordered boron networks, a – B:H, analogous to the glassy carbons. Among the silicides (Table 4), we see a fairly plausible and potentially stable ultrahard compound B_2Si. Table 5, is surprising, in that a host of fairly stable and ultrahard phases are proposed. Among the beryllides (Table 3), the hardest and most plausible projected materials are Be_2H and Be_2Li.

4.1. Observations from network skeletonization

Using a feedforward network trained on the same cumulative database as that used for Creativity Machine Cascade, it was possible to map from the electronic

Table 3
Ultrahard beryllides imagined by the Creativity Machine and ranked by descending chemical stability

A	x	B	y	ΔG_f^0 (kcal mole^{-1})	MP (°C)	Mohs Hardness	Density (g cm^{-3})	xtala(Å)	xtalb(Å)	xtalc(Å)
Be	1.85	Cl	1.72	−177.10	864	9.22	2.96	8.34	8.02	5.76
Be	2.00	Hg	1.00	−63.65	1417	9.31	7.72	5.52	5.57	4.13
Be	2.00	Tl	1.00	−47.00	611	9.53	7.46	5.49	5.49	3.79
Be	1.85	H	1.04	−42.09	929	9.68	1.61	8.61	8.16	4.56
Be	1.85	Li	1.04	−36.39	925	9.66	1.61	8.69	8.26	4.64
Be	2.00	Pb	1.00	−18.95	599	9.69	7.58	7.24	6.89	3.78
Be	1.85	Na	1.04	−16.05	915	9.50	1.74	8.91	8.53	4.92
Be	2.00	Mg	1.00	−8.46	1299	9.35	1.95	5.36	5.68	3.68
Be	2.08	Bi	1.80	−3.13	1235	9.54	7.84	9.75	8.98	3.95
Be	1.85	K	1.04	−0.47	913	9.18	1.99	8.93	8.64	5.14

Table 4
Ultrahard silicides imagined by the Creativity Machine and ranked by descending chemical stability

A	x	B	y	ΔG_f^0 (kcal mole^{-1})	MP (°C)	Mohs Hardness	Density (g cm^{-3})	xtala(Å)	xtalb(Å)	xtalc(Å)
Pd	2.00	Si	1.00	−111.51	1072	9.09	9.13	5.20	5.31	5.59
N	2.00	Si	1.00	−102.71	2796	9.61	2.34	6.23	6.30	4.22
P	2.00	Si	1.00	−89.62	1505	8.85	2.08	6.41	6.45	4.80
B	1.73	Si	1.35	−53.88	2076	9.66	2.81	5.18	5.46	4.01
Si	1.89	Tm	1.11	−33.59	2052	8.24	6.13	7.79	7.74	4.94
Si	1.00	Th	1.00	−31.97	1725	9.12	7.97	7.39	7.78	4.91
Si	1.00	U	1.00	−26.66	1745	9.24	9.45	7.29	7.62	5.10
Si	1.88	Tb	1.12	−25.80	1984	8.48	5.86	7.54	7.63	5.17
Si	1.86	Tl	1.54	−10.31	1872	8.17	8.28	7.26	7.10	3.43
Mn	1.75	Si	1.31	10.62	1166	8.43	6.14	4.54	4.87	5.54

Table 5
Ultrahard intermetallics imagined by the Creativity Machine and ranked by stability

A	*x*	B	*y*	ΔG_f^0 (kcal mole^{-1})	MP (°C)	Mohs Hardness	Density (g cm^{-3})	xtala(Å)	xtalb(Å)	xtalc(Å)
Cm	2.09	Hg	1.37	−88.70	1742	9.93	14.80	4.55	4.68	4.77
Gd	2.09	Hg	1.37	−84.86	634	9.96	11.15	5.54	5.58	5.05
Gd	2.00	Ti	1.00	−80.09	1222	9.96	10.61	5.61	*5.59*	4.64
Nb	2.29	Tl	1.66	−64.78	1137	9.98	10.71	*6.59*	6.56	6.89
Th	2.12	Tl	1.60	−59.63	1176	9.96	12.49	4.47	4.60	4.69
Mo	2.20	Tl	1.64	−53.14	1246	9.99	10.79	6.59	6.50	6.73
Cf	2.05	Ir	1.21	−46.75	2094	9.94	17.39	5.45	5.45	5.64
Ru	2.03	Tl	1.61	−33.13	1117	9.99	10.38	6.50	6.30	6.42
Bk	2.13	Ir	1.27	−32.20	2191	9.98	17.06	5.35	5.35	5.33
Rh	1.96	Tl	1.61	−25.58	668	9.99	10.06	6.36	6.13	6.28
Gd	2.10	Os	1.22	−25.33	2384	9.94	18.03	7.37	7.22	6.76
Mo	2.21	Pb	1.73	−19.83	1310	9.99	9.43	6.42	6.37	7.31
Os	1.89	Po	1.69	−12.60	2491	9.91	12.07	5.86	5.77	6.05
Ag	1.87	Ir	1.38	−11.90	1757	9.92	13.07	4.42	4.43	5.05
Bk	2.13	Re	1.21	−11.61	2893	9.98	18.32	6.25	6.13	6.14
Mo	2.21	Os	1.33	−10.58	2319	9.98	15.15	7.10	6.88	5.31

representation and stoichiometry to 11 nodes representing the respective Mohs scale classes (i.e., 0–1, 1–2,..., 9–10). In Figs. 14 and 15, we see the network skeletonization (see for instance McMillan et al., 1991 [6]) for both the softest and hardest Mohs scale classes, representing the most significant of weights within each of these respective mappings. We note that for the softest compounds, all electronic and stoichiometric parameters figure fairly uniformly in the underlying schema determining that softness. The analogous skeletonization for ultrahard binary compounds shows the predominate role of one constituent element over another, element one, in determining overall hardness. This would indicate that for the most part, within the ultrahards, a single element dominates in producing an already hard skeletal framework that is now occupied by some less crucial sublattice. This would therefore correspond to the well known cases of carbides, borides, beryllides, silicides, to form intrinsically hard sublattices that may be dressed with softer sublattices. We see that the electronic configuration of element one conspires with the stoichiometry of element two, through a hidden layer node, to determine overall lattice hardness. This observation suggests that given sufficiently soft sub-lattice material (ie., element two), the harder lattice may be diluted, thus degrading its effective bulk modulus. Also note the significant role played by the s and p electrons within the harder framework lattice, generally corresponding to the well known bonding schemes within C, B, Be and Si sublattices.

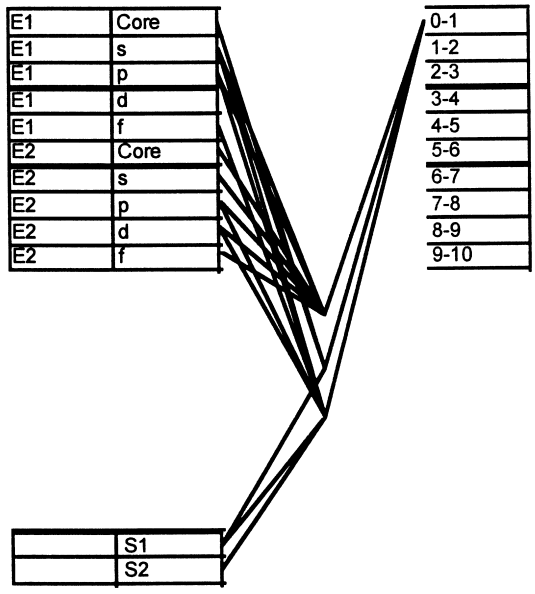

Fig. 14. Network skeletonization for ultrasoft materials.

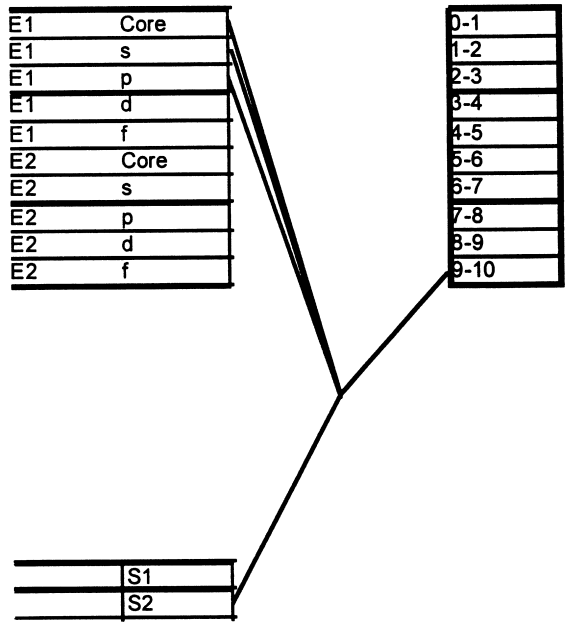

Fig. 15. Network skeletonization for ultrahard materials.

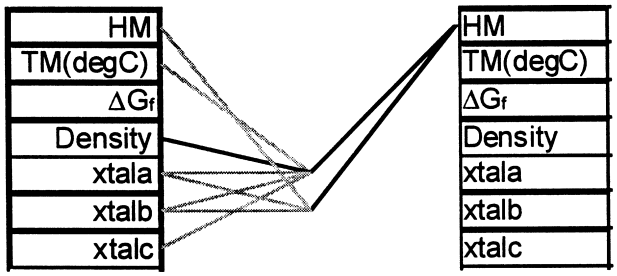

Fig. 16. Network skeletonization of auto-associative net interrelating materials properties. The revealed connection trace shows the prevalent role of melting point, density, and lattice constants in determining Mohs scale hardness. Note that overall, free energy of formation is not generally a strong indicator of hardness. (Black weights are strongest, while gray weights are of secondary importance.)

Training and then skeletonizing (Fig. 16) an auto-associative net trained on the projected properties of 10 000 of the generated compounds reveals the critical factors that generally contribute to hardness. The foremost observation about this schema analysis is that free energy of formation may not be a critical determiner of hardness, as evidenced by the large numbers of ultrahard intermetallics generated by the Creativity Machine.

Further clues as to the underlying hardness trends of binary compounds comes from a plot of Mohs scale hardness versus free energy density for the 4000 hardest materials projected by the Creativity Machine (Fig. 17). Within such a plot we see a thick band of ultrahard materials forming a vertical trunk and showing very little standard free energy stabilization. This class of materials is therefore either, (1) composed of single element phases, having by definition zero free energy of formation, or (2) hetero-nuclear compounds that are inherently unstable from a chemical perspective. This latter case may be representative of a solid made relatively incompressible due to such physical mechanisms as the entrapment of some incompressible atomic core within a chemically inert framework. Examples of such compounds would include the borides such as LaB_6 and its analogs, in which larger cation radius generally results in increased bulk modulus and hardness. We note that the hardest material discovered so far within this class is diamond, with a positive free energy of formation of 0.69 kcal mole^{-1}.

Projecting from this main trunk are multiple classes of binary compounds generally obeying a linear relationship between hardness and chemical free energy per unit volume. The highest of these branches, marked Mo, is dominated by Molybdenum-based intermetallics such as $Mo_{2.21}Pb_{173}$ (Table 5), with an anticipated Mohs hardness of 9.99. A softer, but more pronounced branch occurs at the line marked Pt largely containing Platinum-based intermetallics such as $Pd_{1.77}Pt_{1.41}$. We note that each of these lines generally follow the kind of relationship intimated by Eq. (1), with modulus roughly scaling now as d^{-3} [10]. Another interesting trend line, designated Rh, represents a family of Rhodium compounds that starts at upper right with the intermetallic $Rh_{1.96}Tl_{1.6}$ (See Table 5), with a predicted Mohs scale hardness of 9.99 and terminating with the softer material $Rh_{1.38}Pd_{1.86}$ shows a positive slope, perhaps indicating the presence and dominance of ionic bonding within this structural series. Another potential trend line, labelled X, runs parallel with the Rh series. However, there was no obvious commonality among compounds falling on this line.

Similar analysis (Fig. 18) for the carbides, borides, beryllides, and silicides predicted by the Creativity Machine linear trend lines with significantly more scatter than

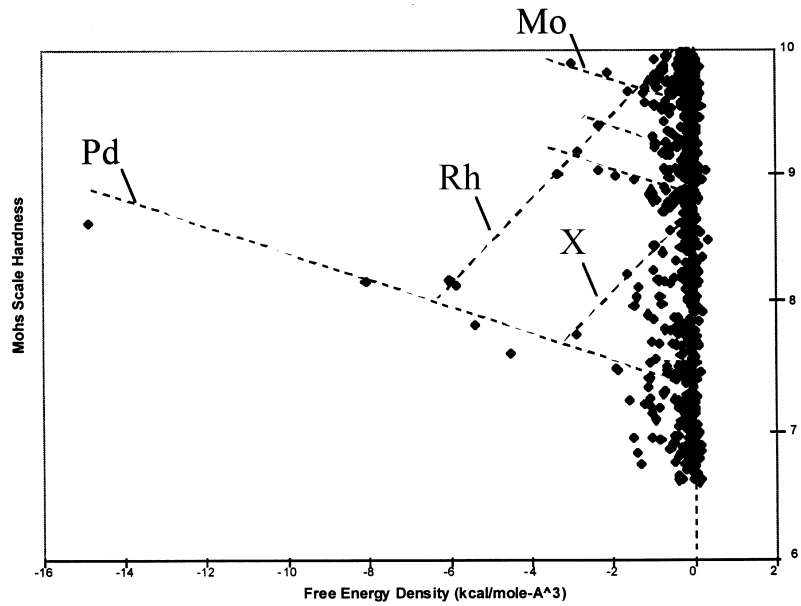

Fig. 17. Plot of Mohs scale hardness versus chemical free energy per unit volume shows pronounced linear branches.

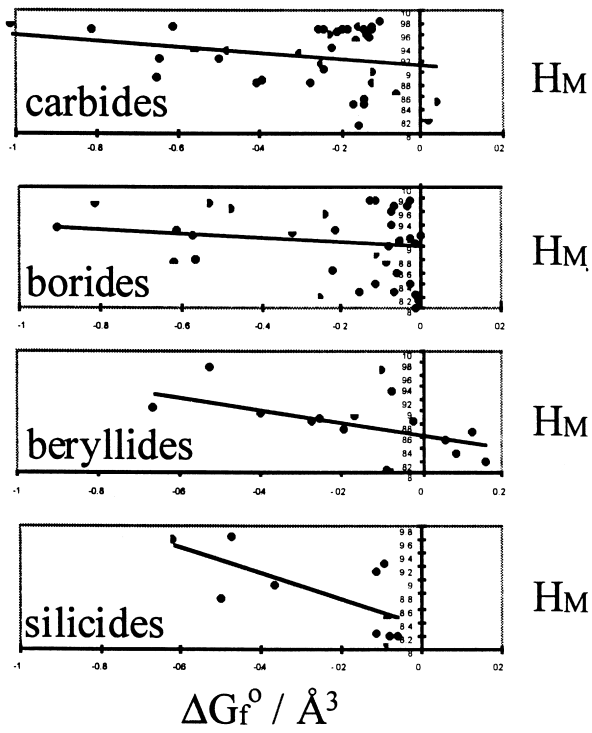

Fig. 18. Linear regression fits to data plots of Mohs scale hardness, H_M, versus approximate free energy density ($\Delta G_f^0/\text{Å}^3$) for ultrahard carbides, borides, beryllides, and silicides discovered by the binary compound Creativity Machine. In each case, we note the roughly linear fit to the data, indicating a dependence of hardness upon the lattice spacings, similar to the pattern of Eq. (1). We also note that each of these ultrahard groupings splits off from the main vertical trunk of Fig. 17, where we anticipate the elemental prototype structure. Proceeding from right to left, we expect an increase in bond energy density, and hence hardness, as covalent bonding between the base lattice and its interstitial lattice develops. For the carbide branch, for instance, the branch appears to originate from two sources, near $H_M = 10$ and $H_M \sim 8$, perhaps representing diamond and graphitic phases respectively. As the carbon lattice is dressed with interstitial atoms, we see an overall increase in hardness in proceeding from right to left. The same trend is evident in the borides, beryllides, and silicides, with the latter grouping manifesting the greatest dependence of hardness upon the free energy density.

the global plot of Fig. 17. Proceeding from right to left within any of these linear series, we begin with a native structure such as a bare carbon or boron cage lattice, typically possessing little if any standard free energy of formation. Interstitially dressing each of these already hard skeletal lattices with a second element generally creates the opportunity for chemical bonding that now incrementally changes hardness through higher bond energy per unit volume. This picture is entirely consistent with the schema analysis of Fig. 15, where we see the dominance of one element in determining overall hardness, typically through elemental networks formed via s and p hybrid orbitals (i.e., the sp^2 and sp^3 hybridization prevalent within carbon, boron, beryllium and silicon lattices).

5. Conclusions

5.1. The methodology

Having ‘seen’ a limited number of chemical compounds, an auto-associative neural network may accurately generalize the rules of stoichiometry between arbitrary elements A and B, allowing us to predict the formulas of as yet unobserved chemical compounds. Making such a network chaotic, we produce a powerful engine that may be used to explore large chemical search spaces, readily identifying binary compounds fulfilling a wide range of desired physical or chemical attributes. Run without specific search criteria, such a chemical Creativity Machine may generate an ever-expanding database of both known and unknown materials along with accompanying estimated physical and chemical attributes. While admittedly sacrificing some accuracy in the prediction of these materials properties, extensive approximate databases of this kind may be used to scope out whole families of materials satisfying niche design requirements.

A certain degree of caution will always be necessary when dealing with such machine-generated databases. Because we are relying totally upon the internal ‘conversation’ among neural networks within a sizable cascade, each of which has learned subtle, hidden patterns by exposure to large materials databases, we are totally at the mercy of their cumulative ‘expertise’ as well as their potential pathologies. Therefore, it would be reasonable to assume that there will be occasional errors and exaggerations that will be beyond the capabilities of human researchers to readily pinpoint. However, we may reasonably recommend this methodology within the context of increasing efforts in the area of combinatorial and theoretical chemistry, where it is crucial to narrow down the most likely candidate materials fulfilling desired technological needs.

Rather than depend entirely upon blind faith in the predictions of these search engines, we may employ any number of explanation facilities to grasp the cascade’s underlying reasoning. We have employed two very powerful techniques to explore the network schema: (1) the ability to supplement what would otherwise be sparse databases with additional materials examples, as in Fig. 17, and (2) the ability to carry out a ‘post-mortem’ skeletonization of the constituent neural nets to identify underlying physical and chemical schema (Figs. 14–17). Analyzing Creativity Machine output and function in these ways serves as a very useful adjunct to materials research.

5.2. Specific materials discoveries

After considerable exercise of the above data mining techniques, we may summarize the major physical and chemical findings as follows:

1. Diamond, with a Mohs scale hardness of 10.0, remains the hardest of the ultrahard phases, with a number of unexpected contenders closely following (see Tables 1–5). Foremost among these theoretical materials are new carbon phases, hydrogen-deficient polymers of carbon, boron, and beryllium, as well as carbon and boron cage structures containing a range of interstitial atoms.
2. The generated materials database reveals a natural dichotomy into two distinct ultrahard groups. In the first of these, we see little standard free energy of formation, indicating either (a) homonuclear lattice structures with strong chemical bonding or (b) heteronuclear lattices that derive their hardness not from chemical bonding, but from purely physical repulsive interaction among atoms, leading to high bulk modulus. This grouping is represented by the vertical trunk of data points within Fig. 17.
3. Carbides, borides, beryllides, silicides, and a family of intermetallics generally segregate into separate hardness families, each represented by diagonal trend lines similar to those in Fig. 17. Within each of these families, hardness may be incrementally increased by the addition of interstitial atoms that tend to covalently bond with the base cage lattice for which the family line is named.
4. Within each of these covalent branches, we see an approximate $1/d^3$ hardness dependence (where d is a lattice constant) generally similar to that of Eq. (1), but now possessing a residual constant term reflecting the hardness of the underlying homonuclear host lattice (i.e., the intersection of each diagonal branch with the vertical trunk of Fig. 17).
5. A totally unexpected family of ultrahard intermetallics is suggested by the trend line Mo, of Fig. 17. Further analysis is planned to confirm whether this prediction is pathological or real.

Hopefully this study illustrates the power of the Creativity Machine Paradigm in generating new discoveries within any arbitrary conceptual space for which numerical historical data exists. This materials problem therefore serves as a template for autonomous discovery within all realms of human endeavor. Because of its completely connectionist design and its ability to propose possibilities beyond its experience, such machines also present a compelling model of mind and of creative human cognition.

Acknowledgements

I thank Dr. Steven LeClair and Dr. Allen Jackson of Wright Laboratories, who were both central to the motivation and editing of this article.

References

[1] A.Y. Liu, M.L. Cohen, Science 245 (1989) 841–842.
[2] A.Y Liu, M.L. Cohen, Phys. Rev. B, 41(15) 1990.
[3] J.C. Phillips, Bonds and Bands in Semiconductors, Academic Press, New York, 1973.
[4] F.H. Pough, A Field Guide to Rocks and Minerals, The Riverside Press, Cambridge, 1960.
[5] D.E. Rumelliart, G.E. Hinton, R.J. Williams, in: D.E. Rumelhart, J.L. McClelland (Eds.), Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol. 1, MIT, Cambridge, 1986.
[6] C. McMillan, M. Mozer, P. Smolensky, Proceedings of IJCN-International Joint Conference on Neural Networks, IEEE, Piscataway, NJ, 1981, pp. 83–88.
[7] M.A. Tamor, K.C. Hass, J. Mater. Res. 5(11) (1990) 2273, 11/90.
[8] S.L. Thaler, 4-2-4 encoder death, Proceedings of the World Congress on Neural Networks 2, Portland, Oregon, 1993, pp. 180–183.
[9] S.L. Thaler, Neural Networks 8(1) (1995) 55–65.
[10] S.L. Thaler, Bull. Am. Phys. Soc. 40(1) (1995) 592.
[11] S.L. Thaler, Neural networks that create and discover, PC AI, May/June, 1996, pp. 16-21.
[12] S.L. Thaler, The spreadsheet implemented 'Creativity Machine' – its construction, function, and current successes, Adaptive Distributive Parallel Computing Symposium, Dayton, Ohio, 1996.
[13] S.L. Thaler, Is neuronal chaos the source of stream of consciousness?, The World Congress on Neural Networks, Lawrence Erlbaum, 1996.
[14] S.L. Thaler, A proposed symbolism for network-implemented discovery processes, The World Congress on Neural Networks, Lawrence Erlbaum, 1996.
[15] S.L. Thaler, U.S. Patent 6,659,666, Device for the Autonomous Generation of Useful Information, 1997.
[16] S.L. Thaler, A quantitative model of seminal cognition: the creativity machine paradigm, Proceedings of the Mind II Conference, Dublin, Ireland, 1997.