

Creative Thinking

- Claude Shannon

Mar. 20, 1955

Up to 100% of the amount of ideas produced, useful good ideas produced by these signals, these are supposed to be arranged in order of increasing ability. At producing ideas, we find a curve something like this. Consider the number of curves produced here - going up to enormous height here.

A very small percentage of the population produces the greatest proportion of the important ideas. This is akin to an idea presented by an English mathematician, Turig<sup>n</sup>, that the human brain is something like a piece of uranium. The human brain, if it is below the critical lap and you shoot one neutron into it, additional more would be produced by impact. It leads to an extremely explosive . . . of the issue, increase the size of the uranium. Turig says this is something like ideas in the human brain. There are some people if you shoot one idea into the brain, you will get a half an idea out. There are other people who are beyond this point at which they produce two ideas for each idea sent in. Those are the people beyond the knee of the curve. I don't want to sound egotistical here, I don't think that I am beyond the knee of this curve and I don't know anyone who is. I do know some people that were. I think, for example, that anyone will agree that Isaac Newton would be well on the top of this curve. When you think that at the age of 25 he had produced enough science, physics and mathematics to make 10 or 20 men famous - he produced binomial theorem, differential and integral calculus, laws of gravitation, laws of motion, decomposition of white light, and so on. Now what is it that shoots one up to this

part of the curve? What are the basic requirements? I think we could set down three things that are fairly necessary for scientific research or for any sort of inventing or mathematics or physics or anything along that line. I don't think a person can get along without any one of these three.

The first one is obvious - training and experience. You don't expect a lawyer, however bright he may be, to give you a new theory of physics these days or mathematics or engineering.

The second thing is a certain amount of intelligence or talent. In other words, <sup>you have</sup> to have an IQ that is fairly high to do good research work. I don't think that there is any good engineer or scientist that can get along on an IQ of 100, which is the average for human beings. In other words, he has to have an IQ higher than that. Everyone in this room is considerably above that. This, we might say, is a matter of environment; intelligence is a matter of heredity.

Those two I don't think are sufficient. I think there is a third constituent here, a third component which is the one that makes an Einstein or an Isaac Newton. For want of a better word, we will call it motivation. In other words, you have to have some kind of a drive, some kind of a desire to find out the answer, a desire to find out what makes things tick. If you don't have that, you may have all the training and intelligence in the world, you don't have questions and you won't just find answers. This is a hard thing to put your finger on. It is a matter of temperament

probably; that is, a matter of probably early training, early childhood experiences, whether you will motivate in the direction of scientific research. I think that at a superficial level, it is blended use of several things. This is not any attempt at a deep analysis at all, but my feeling is that a good scientist has a great deal of what we can call curiosity. I won't go any deeper into it than that. He wants to know the answers. He's just curious how things tick and he wants to know the answers to questions; and if/<sup>he</sup> sees things, he wants to raise questions and he wants to know the answers to those.

Then there's the idea of dissatisfaction. By this I don't mean a pessimistic dissatisfaction of the world - we don't like the way things are - I mean a constructive dissatisfaction. The idea could be expressed in the words, "This is OK, but I think things could be done better. I think there is a neater way to do this. I think things could be improved a little." In other words, there is continually a slight irritation when things don't look quite right; and I think that dissatisfaction in present days is a key driving force in good scientists.

And another thing I'd put down here is the pleasure in seeing net results or methods of arriving at results needed, designs of engineers, equipment, and so on. I get a big bang myself out of proving a theorem. If I've been trying to prove a mathematical theorem for a week or so and I finally find the solution, I get a big bang out of it. And I get a big kick out of seeing a clever way of doing some

engineering problem, a clever design for a circuit which uses a very small amount of equipment and gets apparently a great deal of result out of it. I think so far as motivation is concerned, it is maybe a little like Fats Waller said about swing music - "either you got it or you ain't." If you ain't got it, you probably shouldn't be doing research work if you don't want to know that kind of answer. Although people without this kind of motivation might be very successful in other fields, the research man should probably have an extremely strong drive to want to find out the answers, so strong a drive that he doesn't care whether it is 5 o'clock - he is willing to work all night to find out the answers and all weekend if necessary. Well now, this is all well and good, but supposing a person has these three properties to a sufficient extent to be useful, are there any tricks, any gimmicks that he can apply to thinking that will actually aid in creative work, in getting the answers in research work, in general, in finding answers to problems? I think there are, and I think they can be catalogued to a certain extent. You can make quite a list of them and I think they would be very useful if one did that, so I am going to give a few of them which I have thought up or which people have suggested to me. And I think if one consciously applied these to various problems you had to solve, in many cases you'd find solutions quicker than you would normally or in cases where you might not find it at all. I think that good research workers apply these things unconsciously; that is, they do these things automatically and if they were brought forth into the conscious thinking that here's

a situation where I would try this method of approach that would probably get there faster, although I can't document this statement.

The first one that I might speak of is the idea of simplification. Suppose that you are given a problem to solve, I don't care what kind of a problem - a machine to design, or a physical theory to develop, or a mathematical theorem to prove, or something of that kind - probably a very powerful approach to this is to attempt to eliminate everything from the problem except the essentials; that is, cut it down to size. Almost every problem that you come across is befuddled with all kinds of extraneous data of one sort or another; and if you can bring this problem down into the main issues, you can see more clearly what you're trying to do and perhaps find a solution. Now, in so doing, you may have stripped away the problem that you're after. You may have simplified it to a point that it doesn't even resemble the problem that you started with; but very often if you can solve this simple problem, you can add refinements to the solution of this until you get back to the solution of the one you started with.

A very similar device is seeking similar known problems. I think I could illustrate this schematically in this way. You have a problem <sup>P</sup> here and there is a solution <sup>S</sup> which you do not know yet perhaps over here. If you have experience in the field represented, that you are working in, you may perhaps know of a somewhat similar problem, call it P', which has already been solved and

which has a solution, S'. All you need to do - all you may have to do is to find the analogy from P' here to P and the same analogy from S' to S in order to get back to the solution of the given problem. This is the reason why experience in a field is so important that if you are experienced in a field, you will know thousands of problems that have been solved. Your mental matrix will be filled with P's and S's unconnected here and you can find one which is tolerably close to the P that you are trying to solve and go over to the corresponding S' in order to go back to the S you're after. It seems to be much easier to make two small jumps than the one big jump in any kind of mental thinking.

Another approach for a given problem is to try to restate it in just as many different forms as you can. Change the words. Change the viewpoint. Look at it from every possible angle. After you've done that, you can try to look at it from several angles at the same time and perhaps you can get an insight into the real basic issues of the problem, so that you can correlate the important factors and come out with the solution. It's difficult really to do this, but it is important that you do. If you don't, it is very easy to get into ruts of mental thinking. You start with a problem here and you go around a circle here and if you could only get over to this point, perhaps you would see your way clear; but you can't break loose from certain mental blocks which are holding you in certain ways of looking at a problem. That is the reason why very frequently someone who is quite green to a problem will sometimes

come in and look at it and find the solution like that, while you have been laboring for months over it. You've got set into some ruts here of mental thinking and someone else comes in and sees it from a fresh viewpoint.

Another mental gimmick for aid in research work, I think, is the idea of generalization. This is very powerful in mathematical research. The typical mathematical theory developed in the following way to prove a very isolated, special result, particular theorem - someone always will come along and start generalizing it. He will leave it where it was in two dimensions before he will do it in  $N$  dimensions; or if it was in some kind of algebra, he will work in a general algebraic field; if it was in the field of real numbers, he will change it to a general algebraic field or something of that sort. This is actually quite easy to do if you only remember to do it. If the minute you've found an answer to something, the next thing to do is to ask yourself if you can generalize this any more - can I make the same, make a broader statement which includes more - there, I think, in terms of engineering, the same thing should be kept in mind. As you see, if somebody comes along with a clever way of doing something, one should ask oneself "Can I apply the same principle in more general ways? Can I use this same clever idea represented here to solve a larger class of problems? Is there any place else that I can use this particular thing?"

Next one I might mention is the idea of structural analysis of a problem. Supposing you have your problem here and a solution

here. You may have too big a jump to take. What you can try to do is to break down that jump into a large number of small jumps. If this were a set of mathematical axioms and this were a theorem or conclusion that you were trying to prove, it might be too much for me to try to prove this thing in one fell swoop. But perhaps I can visualize a number of subsidiary theorems or propositions such that if I could prove those, in turn I would eventually arrive at this solution. In other words, I set up some path through this domain with a set of subsidiary solutions, 1, 2, 3, 4, and so on, and attempt to prove this on the basis of that and then this on the basis of these which I have proved until eventually I arrive at the path S. Many proofs in mathematics have been actually found by extremely roundabout processes. A man starts to prove this theorem and he finds that he wanders all over the map. He starts off and proves a good many results which don't seem to be leading anywhere and then eventually ends up by the back door on the solution of the given problem; and very often when that's done, when you've found your solution, it may be very easy to simplify; that is, to see at one stage that you may have short-cutted across here and you could see that you might have short-cutted across there. The same thing is true in design work. If you can design a way of doing something which is obviously clumsy and cumbersome, uses too much equipment; but after you've really got something you can get a grip on, something you can hang on to, you can start cutting out components and seeing some parts were really superfluous. You really didn't need them in the first place.



Now one other thing I would like to bring out which I run across quite frequently in mathematical work is the idea of inversion of the problem. You are trying to obtain the solution S on the basis of the premises P and then you can't do it. Well, turn the problem over supposing that S were the given proposition, the given axioms, or the given numbers in the problem and what you are trying to obtain is P. Just imagine that that were the case. Then you will find that it is relatively easy to solve the problem in that direction. You find a fairly direct route. If so, it's often possible to invert it in small batches. In other words, you've got a path marked out here - there you got relays you sent this way. You can see how to invert these things in small stages and perhaps three or four only difficult steps in the proof.

Now I think the same thing can happen in design work. Sometimes I have had the experience of designing computing machines of various sorts in which I wanted to compute certain numbers out of certain given quantities. This happened to be a machine that played the game of nim and it turned out that it seemed to be quite difficult. It took quite a number of relays to do this particular calculation although it could be done. But then I got the idea that if I inverted the problem, it would have been very easy to do - if the given and required results had been interchanged; and that idea led to a way of doing it which was far simpler than the first design. The way of doing it was doing it by feedback; that is, you start with the required result and run it back until - run it through its value

until it matches the given input. So the machine itself was worked backward putting range S over the numbers until it had the number that you actually had and, at that point, until it reached the number such that P shows you the correct way. Well, now the solution for this philosophy which is probably very boring to most of you. I'd like now to show you this machine which I brought along and go into one or two of the problems which were connected with the design of that because I think they illustrate some of these things I've been talking about.

In order to see this, you'll have to come up around it; so, I wonder whether you will all come up around the table now.