

Disclosure

of things evolutionists don't want you to know

Volume 24 Issue 7

www.ScienceAgainstEvolution.info

April 2020

FRIENDS WITH LONG FACES

This song was inspired by Garth Brooks' song, "Friends in Low Places." ¹

Verse 1:

Blame it all on their schools, their teachers were fools.
The truth, they just didn't hear.
The last ones to know, they were lied to and so
They were the last ones to see things so clear.
They did not realize they'd been told many lies
When they took that class in the fall.
They thought it was true, but
When the class was all through,
They had learned nothing at all!

Chorus:

Now I've got Friends with Long Faces.
They've discovered new facts
And the truth just chases all the lies away.
But they'll be OK.
Darwin was wrong in so many cases.
They used to believe what the truth replaces.
Now I've got Friends with Long Faces.

Verse 2:

"Darwin was wrong," is the point of my song,
But then, I've said that before.
You know that I'm right, so I'll just say goodnight,
And I'll show myself to the door.
Hey, I didn't mean to cause a big scene
When I took an hour and then,
I just demolished that ivory tower
That they're livin' in.

(repeat chorus)

Ending:

I've got Friends with Long Faces.

¹ Hear it performed by Death Valley Dave at <http://www.scienceagainstevolution.info/music/FWLF.mp3>

AUDITORY COGNITION

Alexa, What's the difference between speech and music?

At various times in my life I earned money as a musician, teacher, and engineer. As such, I am interested in sound from entertainment, communication, and technical points of view. So, I was fascinated by a recent article in the journal, *Science*, by Daniela Sammler titled, "Splitting speech and music." She said,

Speech and music are human universals, and people around the world often blend them together into vocal songs. This entwining of the speech and music cognitive domains is a challenge for the auditory cognitive system. How do listeners extract words and melodies from a single sound wave? The split is surmised to start in the signal: Speech and musical sounds are thought to differ in details of their acoustic structure and thus activate different receptive preferences of the left and right auditory cortices of the brain. On page 1043 of this issue, Albouy *et al.* provide evidence for the biophysical basis of the long-debated, yet still unresolved, hemispheric asymmetry of speech and music perception in humans. They show that the left and right auditory regions of the brain contribute differently to the decoding of words and melodies in songs.

Research on the nature of hemispheric asymmetries started in 1861, when French anatomist Pierre Paul Broca astounded his Parisian colleagues with the observation that speech abilities are perturbed after lesions in the left, but not right, brain hemisphere.²

Sammler's introduction to Albouy's article, and Albouy's article itself, both have to do with which portions of the brain do the signal processing. We don't really care about where the signal processing occurs. Our interest has to do with how speech and music are perceived. What is the necessary biological hardware and cranial software that are required for auditory cognition, and is it even remotely plausible that evolution could have produced this capability?

Today, influential neuroacoustic models seek reasons in the specific computational requirements imposed by the structure of speech and musical sounds. For example, proper speech perception hinges strongly (but not solely) on the ability to process short-lived temporal modulations that are decisive for discriminating similar-sounding words, such as "bear" from "pear." By contrast, proper music perception requires, among others, the ability to process the detailed spectral composition of sounds (frequency fluctuations).³

Here's how they conducted the experiment (in their words):

One hundred *a cappella* songs in each language were recorded following a 10×10 matrix with 10 melodies (number code) and 10 sentences (letter code). Stimuli were then filtered either in the spectral or in the temporal dimension with five filter cutoffs, resulting in 1000 degraded stimuli for each language.

We first investigated the importance of STM [SpectroTemporalModulation] rates on sentence or melody recognition scores in a behavioral experiment (Fig. 2A). Native French ($n = 27$) and English ($n = 22$) speakers were presented with pairs of stimuli and asked to discriminate either the speech or the melodic content. Thus, the stimulus set across the two tasks was identical; only the instructions differed.⁴

In plain English, they created 100 songs (without musical accompaniment) by singing 10 different sentences to 10 different melodies. They corrupted these 100 songs 10 different ways by messing with the pitch or timing to create 1000 test songs. They played pairs of these songs to 27 French speakers, and 22 English speakers, and asked them if melodies were the same or different. Then they played the same pairs of songs and asked them if the lyrics were the same or different.

At the same time, they were monitoring the listeners' brains to see what parts of the brain were stimulated when they were trying to determine if the melodies or lyrics were the same.

Years of debate have centered on the theoretically important question of the representation of speech and music in the brain. Here, we take advantage of the STM framework to establish a rigorous demonstration that: (i)

³ *ibid.*

⁴ Albouy, *et al.*, *Science*, 28 Feb. 2020, "Distinct sensitivity to spectrotemporal modulation supports brain asymmetry for speech and melody", pp. 1043-1047, <https://science.sciencemag.org/content/367/6481/1043>

² Daniela Sammler, *Science*, 28 Feb. 2020, "Splitting speech and music", pp. 974-976, <https://science.sciencemag.org/content/367/6481/974>

perception of speech content is most affected by degradation of information in the temporal [time] dimension, whereas perception of melodic content is most affected by degradation in the spectral [frequency/pitch] dimension; (ii) neural decoding of speech and melodic contents primarily depends on neural activity patterns in the left and right AC regions, respectively; (iii) in turn, this neural specialization for each stimulus domain is dependent on the specific sensitivity to STM rates of each auditory region; and (iv) the perceptual effect of temporal or spectral degradation on speech or melodic content is mirrored specifically within each hemispheric auditory region (as revealed by mutual information), thereby demonstrating the brain-behavior relationship necessary to conclude that STM features are processed differentially for each stimulus domain within each hemisphere.⁵

That probably didn't mean much to you—but I spent decades developing signal processing algorithms which analyzed radar returns or visual images to distinguish a target from the background. Depending upon the target of interest and background clutter, these algorithms either worked in the “time domain” or the “frequency domain.”

DOMAINS

Since you might not be familiar with the time domain and frequency domain, here's a short explanation.

A touchtone phone dials numbers in the frequency domain. It determines what button you have pressed by the musical pitch each button makes. Pressing the 3 button produces a different sound than pressing the 8 button. The sound doesn't have to last very long for the telephone system to recognize the frequencies in the sound when dialing the phone.

The old-fashioned rotary phones dialed numbers in the time domain. They produced a series of clicks. (Young readers might not have ever dialed a rotary phone, so you need to know that dialing a 3 produced 3 clicks, and dialing an 8 produced 8 click, and so on.) Each click caused a rotary switch to move one position. When enough time went by without a click, the switchboard would recognize that the sequence had ended for that digit, and would move on to set the rotary switch for the next digit in the phone number. The next series of clicks would set the rotary switch to the next digit in the phone number.

⁵ *ibid.*

If you are old enough to remember dialup modems, you will remember the series of differently pitched beeps the modem made when connecting to AOL. Information was being passed in the frequency domain. If you are really old, and know Morse Code, you know that the information was conveyed as a series of short and long beeps called “dots” and “dashes.” The pitch of the beeps didn't matter. All that mattered was the length of the beeps, and the time between them. Morse Code worked in the time domain.

Albouy and his associates have discovered that the human brain not only uses time domain and frequency domain algorithms to extract information from sound waves, the brain uses different algorithms to extract different kinds of information.

The article ends by saying,

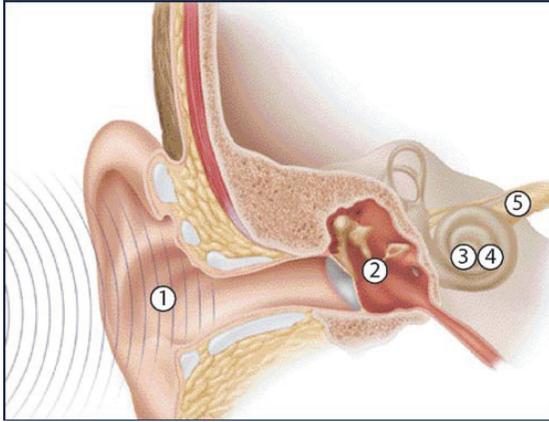
Our study shows that in addition to speech, this theory can be applied to melodic information, a form-bearing dimension of music. Humans have developed two means of auditory communication: speech and music. Our study suggests that these two domains exploit opposite extremes of the spectrottemporal continuum, with a complementary specialization of two parallel neural systems, one in each hemisphere, that maximizes the efficiency of encoding of their respective acoustical features.⁶

We appreciate the fact that they said, “Humans have developed two means ...” rather than “Humans have evolved two means ...”; but we would have preferred them to have said, “Humans have two means ...”. Their research indicates that humans have two means of processing auditory information; but they never claim that auditory communication evolved, perhaps because they understand the complexity. “Developed” (which implies change over time) is as close as they came to crediting evolution for the capability.

That brings us to our discussion of the means of auditory processing, and whether or not it could have evolved. Here are the best explanations we found on the Internet.

⁶ *ibid.*

How the Ear Works



Picture courtesy of Cochlear Ltd.

Here is how the ear works normally:

1. Sound is transmitted as sound waves from the environment. The sound waves are gathered by the outer ear and sent down the ear canal to the eardrum.
2. The sound waves cause the eardrum to vibrate, which sets the three tiny bones in the middle ear into motion.
3. The motion of the bones causes the fluid in the inner ear or cochlea to move.
4. The movement of the inner ear fluid causes the hair cells in the cochlea to bend. The hair cells change the movement into electrical pulses.
5. These electrical impulses are transmitted to the hearing (auditory) nerve and up to the brain, where they are interpreted as sound.⁷

Remarkably, the ‘hair cells’ in the cochlea are tuned to respond to different sounds based on their pitch or frequency of sounds. High-pitched sounds will stimulate ‘hair cells’ in the lower part of the cochlea and low-pitched sounds in the upper part of the cochlea.

What happens next is even more remarkable because, when each ‘hair cell’ detects the pitch or frequency of sound to which it’s tuned to respond, it generates nerve impulses which travel instantaneously along the auditory nerve.

These nerve impulses follow a complicated pathway in the brainstem before arriving at the hearing centres of the brain, the auditory cortex. This is where the streams of nerve impulses are converted into meaningful sound.

All of this happens within a tiny fraction of a second...almost instantaneously after sound waves first enter our ear canals. It is very true to say that, ultimately, we hear with our brain.⁸

Let’s look at the hearing process in more detail, and examine the plausibility of each part evolving by chance.

The funnel-shaped outer ear collects sound from a large area and concentrates it onto the smaller eardrum. It’s plausible that could have happened by accident. That doesn’t mean it did happen by accident—but it could have happened by accident.

Next we come to those three little bones in the middle ear. Because they resemble some bones in the jaws of some reptiles, evolutionists believe that some sort of birth defect caused these bones to develop in the ear instead of in the mouth, and this luckily improved hearing so much that this beneficial mistake established itself in the mammal population. They say,

The evolution of mammalian auditory ossicles [ear bones] is one of the most well-documented and important evolutionary events, demonstrating both numerous transitional forms as well as an excellent example of exaptation, the re-purposing of existing structures during evolution.

In reptiles, the eardrum is connected to the inner ear via a single bone, the stapes or stirrup, while the upper and lower jaws contain several bones not found in mammals. Over the course of the evolution of mammals, one lower and one upper jaw bone (the articular and quadrate) lost their purpose in the jaw joint and were put to new use in the middle ear, connecting to the stapes and forming a chain of three bones (collectively called the ossicles) which amplify sounds and allow more acute hearing. In mammals, these three bones are known as the malleus, incus, and stapes (hammer, anvil, and stirrup respectively).

... The mammalian middle ear contains three tiny bones known as the ossicles: malleus, incus, and stapes. The ossicles are a complex system of levers whose functions include: reducing the amplitude of the vibrations; increasing the amount of energy transmitted.⁹

The idea that jaw bones migrated into the ear, (where they just happened to be useful) is nonsense—but the statement that these bones increase the amount of energy transmitted is absolutely true. They do this by something engineers call, “impedance matching.”

⁷ <https://www.umms.org/ummc/health-services/hearing-balance/patient-information/how-ear-works>

⁸ <https://www.hearinglink.org/your-hearing/about->

[hearing/how-the-ear-works/](https://www.hearinglink.org/your-hearing/about-hearing/how-the-ear-works/)

http://en.wikipedia.org/wiki/Evolution_of_mammalian_auditory_ossicles

IMPEDANCE MATCHING

Impedance is resistance to motion.

If you ever rode a 10-speed bicycle, you have had some personal experience with impedance matching. When riding uphill, you used a lower gear. When riding on a level surface with the wind at your back, you used the highest gear.

In low gear, it is easy to pedal, but you have to pedal more times to go a certain distance. High gear moves you farther down the road with each revolution, but it is harder to push the pedals. Low gear presents the lowest impedance to your legs.

By changing gears, you change the impedance experienced by your leg muscles. You naturally try to select the impedance that optimizes energy transfer.

The physical definition of “work” is “force times distance.” To do a certain amount of work, you either have to apply lots of force for a short distance, or a little bit of force for a long distance. Gears and levers allow you to trade force for distance (and vice versa) to optimize energy transfer. If you have ever jacked up a car to change a tire, you used pounds of force to move the jack handle a few feet to lift tons of weight a few inches.

So what does all this talk about gears and levers, force, distance, and impedance have to do with hearing? Impedance matching is necessary for optimum transfer of sound energy to the inner ear. We will get to that in a minute; but let’s go to the swimming pool first.

REFLECTION CAUSED BY IMPEDANCE MISMATCH

No doubt at some point in your life you stood in the shallow end of a swimming pool and heard all the people talking and kids yelling. Then you dipped your head under water and could not hear the talking very well. Instead, you heard all sorts of underwater noises. Why could you not hear the people talking when you put your head under water? The answer is, “impedance mismatch.”

Air is a low impedance transmission media. It is easy to move air molecules back and forth, so they have to be moved a long distance to transmit energy. Water is a high impedance transmission media. Water molecules are heavy, and don’t need to move very far to transmit energy.

When a sound wave passing through the air hits the surface of the water, the air molecules are too light to push the heavy water molecules, so the sound wave just bounces off the water. When a sound wave traveling through the water reaches the surface, the water molecules move so little that they don’t move the air molecules enough to

be heard. The sound wave can’t get out of the water and into the air, so it reflects back into the water.

Whenever there is an impedance mismatch, energy is reflected rather than transferred.

The fluid inside your ear is high impedance. The bones in your ear are little levers that trade distance for force. They convert the low impedance sound waves in the air to high impedance sound waves in the fluid in your ear.

Engineers know how important it is to achieve the proper impedance matching to facilitate energy transfer. To think that bones that happened to grow in the wrong place (that is, in the ear instead of the jaw) would just happen to have the proper gear ratio to efficiently transfer sound waves from air to liquid, is just foolish. “But then, I’ve said that before.” 😊¹⁰ Let’s move on to what we haven’t said before.

SIGNAL PROCESSING

Once the sound energy has been transferred to the inner ear, many hair-like sensors detect the sound. Each sensor responds to a different frequency (musical pitch). From an engineering perspective, many detectors working in parallel send the data to the brain to process the sound in the frequency domain.

When listening to music, the brain recognizes the sequences of different pitches which form the melody. Rhythm (timing) is important, too—but I think it is easier to recognize most songs from the melody rather than the rhythm. That’s just my opinion, but it seems to be born out by Albuoy’s research.

SPEECH, SPEECH!

Albuoy’s research indicates the brain processes speech differently from music, and claims that a different part of the brain processes speech. Regardless of where the processing happens, I believe that speech recognition is more difficult than music recognition, based on my experience with pattern recognition algorithms.

Younger readers, who grew up with Siri and Alexa, might take speech recognition for granted. They don’t realize that engineers worked on the speech recognition problem for years, and made progress very slowly. Early speech-to-text programs had to be trained by the user, who repeated a list of key words over and over again, until the program recognized those words (and only those words) correctly. Vocabulary was limited to those few words. It seemed like

¹⁰ Disclosure, April 2012, “I Heard it Through My Jaw Bones”, <http://scienceagainstevolution.info/v16i7f.htm>

speaker-independent voice recognition (in which it didn't matter whose voice it was) was an impossible dream.

We want to describe speech recognition in sufficient detail to accurately convey how difficult it is, without getting too technical to understand. Since Ed Grabianowski did an excellent job of this, here are a few quotes from his description.

The [auditory] system filters the digitized sound to remove unwanted noise, and sometimes to separate it into different bands of frequency (frequency is the wavelength of the sound waves, heard by humans as differences in pitch). It also normalizes the sound, or adjusts it to a constant volume level. It may also have to be temporally aligned. People don't always speak at the same speed, so the sound must be adjusted to match the speed of the template sound samples already stored in the system's memory.¹¹

In other words, there needs to be some filtering to separate the speech from the background noise. Then it has to be "temporally (not temporarily) aligned" with known sounds. "Temporal alignment" means to compare sound waves point by point in time. The speech might have to be sped up, or slowed down, so that the speech is the same length as the pattern it is being compared to. This isn't hard to do—if you know what you are trying to align it to; but if you knew that in advance you wouldn't need to align it (because you already knew what it was). You have to speed the speech up, and slow it down by different amounts, many times to see if it aligns with lots of different patterns.

Next the signal is divided into small segments as short as a few hundredths of a second, or even thousandths in the case of plosive consonant sounds -- consonant stops produced by obstructing airflow in the vocal tract -- like "p" or "t." The program then matches these segments to known phonemes in the appropriate language. A phoneme is the smallest element of a language -- a representation of the sounds we make and put together to form meaningful expressions. There are roughly 40 phonemes in the English language (different linguists have different opinions on the exact number), while other languages have more or fewer phonemes.¹²

Apparently linguists are still arguing about

exactly how many phonemes there are. One might think that since there are 26 letters in the alphabet, there are 26 different sounds. But, you know that English is difficult because that isn't the case. (Do you have 26 phonemes? Go phish! ☺)

Text-to-speech is easier than speech-to-text, so it was developed first; but even that was hard in the 20th century. The mechanical voices on some early GPS systems really butchered some street names. In the beginning, it was really hard to make robotic voices sound human. Older readers will know I speak the truth!

The 40 (or so) phonemes that were developed for text-to-speech have now become the basis for speech recognition. The speech-recognition algorithm breaks speech into phonemes.

The next step seems simple, but it is actually the most difficult to accomplish and is the focus of most speech recognition research. The program examines phonemes in the context of the other phonemes around them. It runs the contextual phoneme plot through a complex statistical model and compares them to a large library of known words, phrases and sentences. The program then determines what the user was probably saying and either outputs it as text or issues a computer command.¹³

Clearly, to build a computer that converts speech to text takes lots of memory, and lots of processing power. So, the human brain requires a comparable amount of memory and processing to achieve the same result. But there's more!

Accents, dialects and mannerisms can vastly change the way certain words or phrases are spoken. Imagine someone from Boston saying the word "barn." He wouldn't pronounce the "r" at all, and the word comes out rhyming with "John." Or consider the sentence, "I'm going to see the ocean." Most people don't enunciate their words very carefully. The result might come out as "I'm goin' da see tha ocean." They run several of the words together with no noticeable break, such as "I'm goin'" and "the ocean." Rules-based systems were unsuccessful because they couldn't handle these variations. This also explains why earlier systems could not handle continuous speech -- you had to speak each word separately, with a brief pause in between them.¹⁴

Russian speech is harder for me because all the words seem to run together!

¹¹ Ed Grabianowski, "How Speech Recognition Works", <https://electronics.howstuffworks.com/gadgets/high-tech-gadgets/speech-recognition1.htm>

¹² *ibid.*

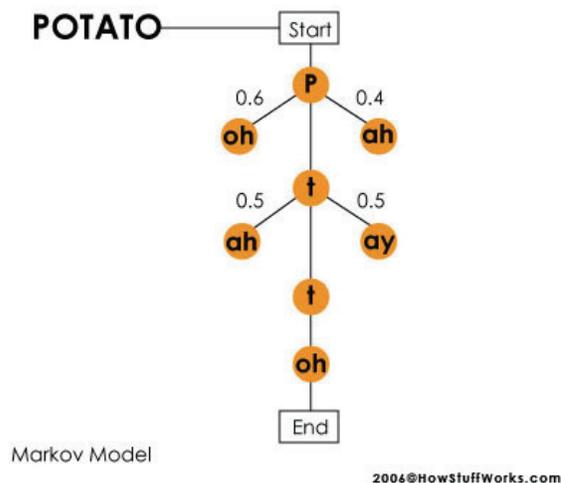
¹³ *ibid.*

¹⁴ Ed Grabianowski, "Speech Recognition and Statistical Modeling", <https://electronics.howstuffworks.com/gadgets/high-tech-gadgets/speech-recognition2.htm>

THE MARKOV MODEL

The Markov Model was developed because rules-based systems were unsuccessful.

How Speech Recognition Works



In this model, each phoneme is like a link in a chain, and the completed chain is a word. However, the chain branches off in different directions as the program attempts to match the digital sound with the phoneme that's most likely to come next. During this process, the program assigns a probability score to each phoneme, based on its built-in dictionary and user training.

This process is even more complicated for phrases and sentences -- the system has to figure out where each word stops and starts. The classic example is the phrase "recognize speech," which sounds a lot like "wreck a nice beach" when you say it very quickly. The program has to analyze the phonemes using the phrase that came before it in order to get it right. Here's a breakdown of the two phrases:

r eh k ao g n ay z s p iy ch

"recognize speech"

r eh k ay n ay s b iy ch

"wreck a nice beach"

Why is this so complicated? If a program has a vocabulary of 60,000 words (common in today's programs), a sequence of three words could be any of 216 trillion possibilities. Obviously, even the most powerful computer can't search through all of them without some help.¹⁵

All languages (including computer languages) have defined syntax and semantics. Syntax is the set of grammatical rules the language uses. Semantics has to do with the meaning of words and symbols. You would realize that the correct

phrase is "recognize speech" rather than "wreck a nice beach" because of semantics. "Recognize speech" makes sense; but "Wreck a nice beach" would make sense only in the context of an environment catastrophe.

The grammar checker of Microsoft Word tells me there is an error in this sentence because **its** wrong. "Its" is a possessive pronoun, like "his." The syntax requires that phrase to have a subject and a verb in that clause. The grammar checker of Microsoft Word flags that sentence because **it's** wrong. Its syntax is wrong.

Our brains use syntax and semantics to extract meaning from sounds.

Homonyms cause a problem for speech recognition, too. (Not "to" or "two" because those two words are not syntactically or semantically correct in that sentence.)

Homonyms are two words that are spelled differently and have different meanings but sound the same. "There" and "their," "air" and "heir," "be" and "bee" are all examples. There is no way for a speech recognition program to tell the difference between these words based on sound alone. However, extensive training of systems and statistical models that take into account word context have greatly improved their performance.¹⁶

SIRI, IS EVOLUTION PLAUSIBLE?

When you think carefully about the hearing process in detail, you must appreciate how complex it really is. The eardrum and bones in the middle ear act as levers to match the impedance of the air to the impedance of the fluid in the cochlea (inner ear) for maximum energy transfer. The hairs in the cochlea are sensitive to different frequencies, converting the sound to signals in the frequency domain which are transmitted as electrical impulses to the brain. The brain uses dual processors, one operating in the frequency domain (apparently in one part of the brain), and the other one in the time domain (apparently in another part of the brain, if Albouy, *et al.*, are correct), to extract music and speech from the sound waves. Then it does pattern matching to extract meaning from the speech (or music, in the case of a police car's siren song).

To think that this all happened by evolution is equivalent to thinking that Siri and Alexa just happened to be able to understand your commands and questions by chance.

¹⁶ Ed Grabianowski, "Speech Recognition: Weaknesses and Flaws",

<https://electronics.howstuffworks.com/gadgets/high-tech-gadgets/speech-recognition3.htm>

¹⁵ *ibid.*

APOLOGETICS PRESS

<http://www.apologeticspress.org/>

Providing scripturally sound and scientifically accurate materials in the field of Christian apologetics

While searching for an article for the April 2020 newsletter for our website, I happened upon a link to an article found on the website of the Apologetics Press. Since there is a great deal of information about Creation vs. Evolution on this site, I thought it would be appropriate to describe in some detail what the reader may find by exploring this site.

The main page of the site serves as the launching point to learn about new and archived information of the Apologetics Press. Here you will find an RSS link that will enable you to easily subscribe to the latest news from Apologetics Press. If you are not familiar with what an RSS feed is, "Unlike a subscription email, an RSS feed doesn't require you to use an email account to stay up to date with news and updates on your favorite content." All you must do is use a modern web browser and configure it to support RSS feeds. If you want, you can also just sign up for the E-mail Newsletter.

The main page of the site also provides links to four recent and timely articles that you can explore as well as links to 1) Featured Video, 2) Featured Audio, 3) Featured Article, 4) *Reason & Revelation Journal* Issue and Archive and 5) *Discovery Magazine* Issue and Archive. *Reason & Revelation* is a monthly journal on Christian Evidences for Adults and has been published since 1981. *Discovery* is a monthly Christian evidences magazine for kids and has been published since 1990.

On the main page of the Apologetics Press you will also find a link to a section called, "Creation vs. Evolution." By selecting this link, you arrive at a list of articles providing the article Title, Author and Publishing Date. A search box is provided to allow you to Search By Keyword or Search By Bible Verse. Also, to make it easier to find articles of interest, the articles can be viewed by selecting one of 31 different Topics.

Besides exploring the many different articles about creation and evolution on this website, you can also explore the PDF Books link found on the main page of the site. Here you learn that "Apologetics Press also publishes numerous books, tracts, research articles, article reprints, video tapes, audio tapes, and other items...However, because our work is strictly non-profit, over the next few years (as time and funds allow) we plan to place on our Web site electronic versions of almost all of our books. As additional titles are placed in electronic format, we will add them to the 'PDF Books' section of our Web Site."

As mentioned above, there is much to explore on this website. Just search for topics on the Creation vs. Evolution link or read books on the PDF Books link that you may find of interest.



**You are permitted (even encouraged)
to copy and distribute this newsletter.**

Disclosure, the Science Against Evolution newsletter, is edited by R. David Pogge.

All back issues are on-line at **ScienceAgainstEvolution.info**.