Journal: www.jogg.info
Originally Published: Volume 5, Number 1 (Spring 2009)
Reference Number: 51.003

ON THE PROPAGATION OF MITOCHONDRIAL MUTATIONS

Author(s): Ian Logan

On the Propagation of Mitochondrial Mutations

Ian Logan

Abstract

How does a new mutation get to appear in a mitochondrial DNA sequence? In this paper some ideas as to how mitochondrial mutations appear are discussed and an accompanying Computer Model attempts to demonstrate the feasibility of these ideas. A new mutation is considered to occur in the first instance as an *initial mutational event*. Next it may become the dominant form in the mitochondrial DNA molecules of a mitochondrial nucleoid, then of a single mitochondrion, and finally the mutation may become present in all the cells of an individual where it can be detected by sequencing of the mitochondrial DNA. The source code of the Computer Model is written in Javascript and is reproduced in full. The Model can be run using a web browser.

Introduction

How does a new mutation get to appear in a mitochondrial DNA sequence? This paper attempts to show how the behaviour of mitochondrial mutations might lead to the emergence of a new mutation in a subject's mitochondrial DNA (mtDNA) sequence. Accompanying the paper is a Computer Model that demonstrates the steps involved.

Although the mechanisms of the production and propagation of chromosomal mutations can be explained, the behaviour of mitochondria and the propagation of mutations in the mtDNA are still very poorly understood. This paper, in no way, manages to resolve all the issues, but the steps involved in going from an *initial mutational event* to the final state of *homoplasmy* are discussed.

Initial mutational event	The actual event when a base changes, i.e. a mutation occurs
de novo mutation	A term used by many researchers when referring to a first appearance of a new mutation in an individual's mtDNA sequence (e.g., as a result of an mtDNA test)
Homoplasmy	A state where all mtDNA molecules in an individual contain the same new mutation

Address for correspondence: Ian Logan, ianlogan22@btinternet.com

Received: August 29, 2008; accepted February 5, 2009

Background

The study of mitochondrial DNA in Genetic Genealogy has developed rapidly over the last few years; and stems mainly from the publication of three important papers. The first in 1981 from Anderson, et al. (Anderson, 1981) from the University of Cambridge suggested the sequence of bases that make up the human mtDNA. The second paper from Cann, et al. (Cann, 1987) from the University of California showed that by studying the mutations found in the mtDNA from different individuals it was possible to develop a phylogenetic tree. This paper led directly to the development of the "Out of Africa" theory that uses the pattern of mutations found in persons from different racial groups to show that all of mankind has descended from a single individual who lived in Africa about 200,000 years ago. The third paper from Andrews, et al. produced in 1999 (Andrews, 1999) corrected some errors in the previous work from 1981 and produced a revised version of the Cambridge Reference Sequence (CRS); thereby giving an accurate sequence against which all other mtDNA sequences can be compared.

The Phylogenetic Tree

The phylogenetic tree for human mtDNA is now very extensive as it is found that only persons who have a maternal relationship within a genealogical time frame have identical mitochondrial sequences. And, as one compares the mtDNA sequences from persons who are more and more distantly related there is an increasing number of mutational differences between the sequences.

There are several papers that describe the phylogenetic tree; in particular, the paper from Kivisild, et al. (2006) has some very useful diagrams. This paper also gives

estimates as to when the different branches of the phylogenetic tree were formed. As an example, the diagrams show that Haplogroup H and V have a common source from about 16,000 +/- 6,000 years ago.

The author's own mtDNA sequence shows his maternal origin is in Haplogroup V and has this set of mutational differences from the CRS:

T72C A263G 309.1C 315.1C A750G
A1438G A2706G A2880G G4580A A4769G
C7028T A8860G A15326G C15904T A16162G
T16298C

The list contains 16 mutations of which 14 are base substitutions and 2 are insertions.

As Haplogroup V and Haplogroup H arose from a common source it is possible to split this list of mutations into two parts.

The first part contains the 6 mutations that have occurred after the split from the common source on the line leading to Haplogroup V:

T72C A2880G G4580A C15904T A16162G
T16298C

whilst the second part contains the other 10 mutations that occur on the line to the CRS in Haplogroup H.

A263G 309.1C 315.1C A750G A1438G
A2706G A4769G C7028T A8860G A15326G

However, if the list of mutations as given above is viewed from the CRS, it is necessary to *invert* the mutations in order to show the mutations which occurred after the split from the common source:

G263A C309d C315d G750A G1438A
G2706A G4769A T7028C G8860A G15326A

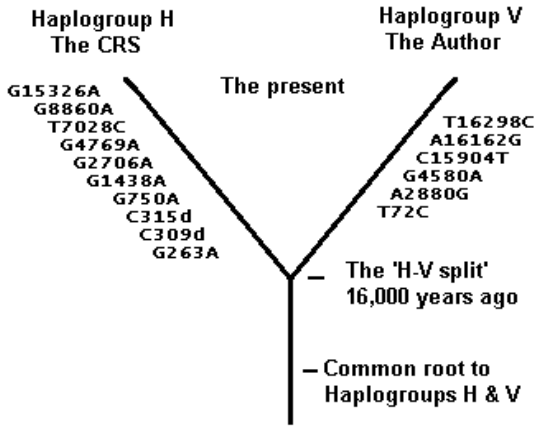

Figure 1. The mutations that have occurred along the lines leading to Haplogroups H and V from HV

Figure 1 illustrates the list of mutations divided into these two parts.

So, over the last 16,000 years these 16 mutations have appeared in the two lines to Haplogroups H and V; but how might this have happened ?

Personification

A point to be made at this stage is that in many places in this paper, and in the accompanying computer model, it has been found useful to *personify* a mutation as it passes through the different stages of its propagation; but it is important to appreciate that all the steps occur at random and a mutation does not have *purpose*, *ambition*, nor *any form of intelligence*.

Modelling the Propagation of Mitochondrial Mutations

A computer model that simulates the propagation of mitochondrial mutations has been written to accompany this paper. This model shows how a new mutation might go through the phases of propagation of a mutation from an *initial mutational event* to the states of heteroplasmy and homoplasmy.

The computer model is available in text form as supplementary file at:

http://www.jogg.info/51/nucloid_model.txt

And as a web page, accompanied by instructions for use and a discussion of the modeling process at:

http://jogg/info/51/nucleoid_model.html

The computer model is set up as a web page and can be viewed with any web browser. However, the model runs appreciably faster under *Google Chrome*, as compared to other browsers, as this browser uses compiled Javascript.

Discussion:

Mitochondria are small organelles found in all the cells of the body. They are mainly involved in the production of "energy" within the cells and have been considered to act as small "power stations." But, whereas most processes within a cell are controlled by the nucleus, mitochondria act independently and are thought to have arisen from bacteria that have formed a symbiotic relationship with the cell. For this reason mitochondria appear to behave more like bacteria than nucleated cells.

Most cells contain relatively few mitochondria, perhaps up to several hundred. But in a fully developed egg cell, an oocyte, there may as many as 200,000. Mitochondria appear to replicate by division and this process is

independent of cell division. So, in a cell that does not divide as it matures, such as an oocyte, the number of mitochondria may become very high. However, in the cells of a tissue where there is a high rate of cell division the number of mitochondria per cell falls to a very low number.

The recent papers by Shoubridge and Wai (Shoubridge, 2007, Wai, 2008) deal particularly well with the subject of the behaviour of mitochondria in mammalian cells and discuss some of the problems concerning the emergence of mitochondrial mutations.

The Propagation of a Mitochondrial Mutation

A new mitochondrial mutation cannot just appear suddenly in every one of an individual's billions of mitochondrial DNA molecules as the biology of mitochondria just does not permit this to happen.

But because mitochondria act like bacteria it is possible to infer that a mitochondrial mutation propagates in the same manner as a mutated bacterial strain is known to arise. In bacteria for the emergence of a new bacterial strain there needs in the first instance to be a mutation occurring in the genetic code of just one organism, this mutated bacterium then goes through many generations until bacteria with this mutation form a new strain. In bacteria a mutation can develop because it offers an advantage to the bacteria, or at random with no advantage being present.

As described with bacteria, the spread of a mitochondrial mutation has to begin with an *initial mutational event* occurring in just one mtDNA molecule and this is followed by the mutation gradually coming to be present in more and more mtDNA molecules, until finally the state of *homoplasmy* is reached when it is considered that every mtDNA molecule in the body has the mutation.

In this paper the gradual increase in the prevalence of a new mutation is described as the propagation of a mutation; and this process which can be split into two phases.

The first phase includes the steps that may be involved in going from an *initial mutational event* through to when many, or all, of the mtDNA molecules, firstly in a single mitochondrial nucleoid, and later a single mitochondrion, have the new mutation. This first phase can occur in any cell in the body, and not necessarily in a cell line that is going to be passed to a descendant.

However, the second phase deals with the special case where the new mutation has present in a germ cell line that goes on to produce the egg cells, the oocytes. In this phase a new mutation propagates and leads to *heteroplasmy* and *homoplasmy* in descendants.

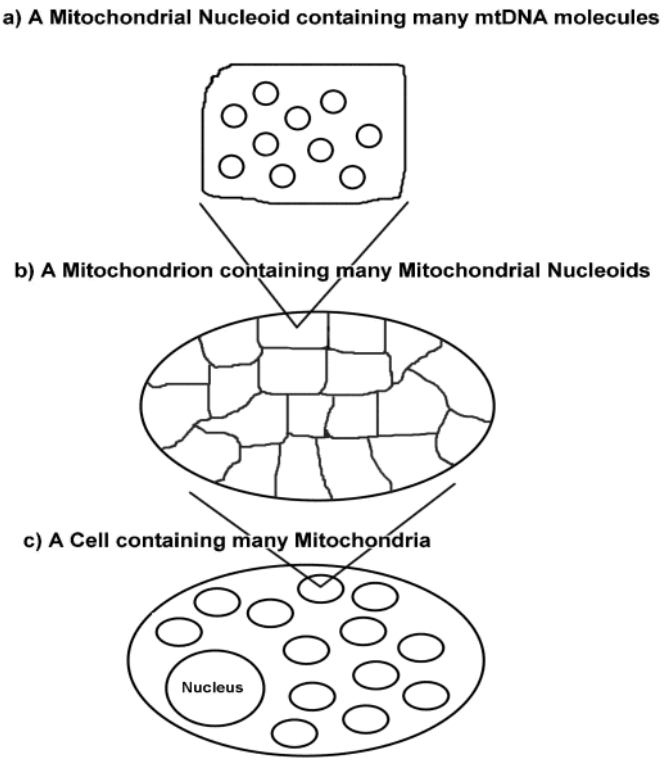

Figure 2. (a) mtDNA molecules in a nucleoid, (b) Nucleoids in a mitochondrion, (c) Mitochondria in a cell.

Figure 2 illustrates the physical relationship between mtDNA molecules, a mitochondrial nucleoid, a mitochondrion and a cell.

Initial Mutational Events

Any mtDNA mutation must in the first instance come from an *initial mutational event* affecting a strand of mtDNA;
and will usually involve either:
- a substitution of one base for another,
- a deletion of one or more neighbouring bases,
or
- the insertion of one of more extra bases,

It is also possible that a *initial mutational event* may, on occasions, be more complex, and lead to several different changes in a mtDNA molecule occurring at the same time. For example, a deletion of bases at one point and their insertion at another may occur.

Presumably, these *initial mutational events* occur as errors when a mtDNA strand is being replicated and are very common. But what is unusual is for an initial event to lead to the development of a new homoplasmic mutation.

In this paper it is assumed that *initial events* occur at random and do not exert any advantage, or disadvantage, to the mitochondrion, cell, or person, in which they are found. In this respect, the discussion of any possible positive or negative *selection* of individuals with certain mtDNA mutations is beyond the scope of this paper.

However, a very interesting paper on *selection* has recently been published by Doublet, et al.,(Doublet, 2008) and whilst the authors are discussing the mitochondria in a crustacean, some of their conclusions on the general behaviour of mitochondria are similar to those expressed in this paper.

Modelling initial mutational events

In the computer model the user is asked to select the number of *initial events* from the range 1 - 10,000,000; and this number then becomes the number of 'runs' that the model will use. In most instances an *initial mutational event* will not propagate, but each mutational event is considered to have the same chance of successful propagation as any other.

Propagation to a Mitochondrial Nucleoid

After its creation as an *initial event*, the next stage in the successful propagation of a mutation is when, by random, many, or all, of the mtDNA molecules in a particular mitochondrial nucleoid have the mutation. In this paper this successful propagation is described as the *capture* of a nucleoid by a mutation; and whereas *initial events* are likely to be very common, a *capture* of a nucleoid is presumably very uncommon.

Unfortunately, the biology of mitochondrial nucleoids is poorly understood - but a nucleoid appears to be a functionally separate area containing just a fraction of the mtDNA molecules found within a single mitochondrion. The recent paper by Gilkerson, et al. (2008) contains a detailed account of what is known about mitochondrial nucleoids, together with a very useful set of references.

An inference that can be drawn from the concept of the *capture* of a mitochondrial nucleoid is that different nucleoids in the same mitochondrion may become captured by different mutations. That is, the mtDNA molecules in one nucleoid may contain one particular mutation whilst other nucleoids may have other mutations

Phase A: Modelling the capture of a mitochondrial nucleoid

In the computer model the user is asked to determine the *Nucleoid Size* from the range 5 - 30, and this value represents the average number of mtDNA molecules in each nucleoid within a mitochondrion.

With the selections of *initial events* and *Nucleoid Size* made, the user can now click A:START to run the model.

Two Sets of Example Inputs and Outputs

Set a --
Number of Mutations = 1000
Nucleoid Size = 10

Mutation 859 captured a nucleoid.

Set b ---
Number of Mutations = 1000000
Nucleoid Size = 20

No nucleoid captured.

These results suggest that with a *Nucleoid Size* of 10 mtDNA molecules per nucleoid, *capture* of a nucleoid may occur, but as the *Nucleoid Size* is increased then nucleoid *capture* becomes very much less likely.

The model will then for each *initial event* in turn show if a new mutation can *capture* a nucleoid.

In simple terms the model replicates a mtDNA molecule chosen at random and then deletes a mtDNA molecule, again chosen at random, over and over, until either the new mutation is lost or it *captures* the nucleoid.

Propagation to a Mitochondrion

Once a new mutation has become found in most, or all, of the mtDNA molecules in a nucleoid, the mutation may then become dominant within a mitochondrion by replication occurring at random, and without suggesting that there might be any selective advantage to the 'mutated' nucleoid.

It is not known how nucleoids are replicated - but they do act as functionally separate organelles. As such it is likely that a nucleoid increases in size as its mtDNA molecules are duplicated; and then the nucleoid divides into two forming daughter nucleoids which each contain roughly half the original number of mtDNA molecules.

This replication of nucleoids appears to happen at random, and this may lead over a number of replication cycles to the situation where all the nucleoids in a mitochondrion are the descendants of just one nucleoid. So, if this nucleoid does contain a new mutation in its DNA molecules the situation will have been reached when all the mtDNA molecules in a mitochondrion have the new mutation. In this paper this is described as the mutation making a successful *capture* of the mitochondrion.

Phase A: Modelling the capture of a mitochondrion

In the computer model the user is asked to determine the *Nucleoid Number* from the range 5 - 20, and this value represents the average number of mtDNA molecules in a mitochondrion. The user also needs to specify 'mitochondrion' as opposed to 'nucleoid'.

Two Sets of Example Inputs and Outputs

Set a ---
Number of Mutations = 1000000
Nucleoid Size = 10
Nucleoid Number = 10

Mutation 113044 captured a mitochondrion.

Set b ---
Number of Mutations = 1000000
Nucleoid Size = 10
Nucleoid Number = 15

No mitochondrion captured.

With the selections of *initial events, Nucleoid Size, Nucleoid Number* and 'mitochondrion' made, the user can now click A:START to run the model.

The model will then for each *initial event* in turn show if a new mutation can *capture* a mitochondrion.

In simple terms the model will as before first try to *capture* a nucleoid, and when successful, proceeds with this nucleoid to see if, at random, it *captures* the mitochondrion.

These results suggest that with a *Nucleoid Number* of 10 mtDNA molecules per mitochondrion *capture* of a mitochondrion may occur, but as the *Nucleoid Number* is increased then mitochondrial *capture* becomes very much less likely.

The Persistence of a Mutated Mitochondrion

Once a mitochondrion has most, or all, of its mtDNA molecules containing a new mutation, it is possible that a fairly stable situation has been reached. Mitochondria are organelles that appear to have a long life within cells so it is likely that a mutated mitochondrion once it has arisen may persist in tissues.

It is interesting to speculate that the phenomenon whereby many new mutations appear in mtDNA taken from cancerous tissue that the mutations have arisen because of the proliferation of persistent mutated mitochondria in the tissue; and not as a result of the actual cancer causing mutations in the mtDNA of its cells.

The Importance of the Germ Cell Line

The previous discussion has concerned itself with how a new mutation might appear in any line of cells in an organism, but in genealogy we are concerned with how mutations are transferred from a person in one generation to their offspring in the next generation. And, for this to happen it is necessary to consider that *mutated mitochondria* are present in the *germ cell line*—meaning that for a new mutation to appear, *mutated mitochondria* need to be present in the egg cells, the *oocytes*.

After fertilisation, a human oocyte enlarges and divides repeatedly to form the first stage of embryonic growth, the *blastula*. This has the form of a hollow ball of cells; and each of the different cells of the *blastula* goes on to divide and make a different cell line; and these cell lines in turn form the different tissues of the body.

One of these lines in a female embryo is the cell line that form the *primordial germ cells* which over time form the next set of *oocytes*; and mitochondria containing mtDNA molecules with a new mutation must be present in this cell line if the new mutation is to be taken forward to the next generation.

Another important cell line to be considered is the one which goes on to produce *buccal* cells (cells forming the inner surface of the cheeks); as sequencing of the mitochondrial DNA is commonly performed on *buccal* swab samples.

But, as before whenever there is replication of cells and their organelles, there is a random allocation of the organelles into the cells. So when different cell lines are produced from the *blastula* it is possible that some cell lines may have a below average number of *mutated mitochondria* in their cells and other cell lines an above average number. If the new *primordial germ cell line* has a higher proportion of *mutated mitochondria* in its cells than were present in the *germ cell line* of the parent then the new mutation has successfully achieved a higher level of propagation.

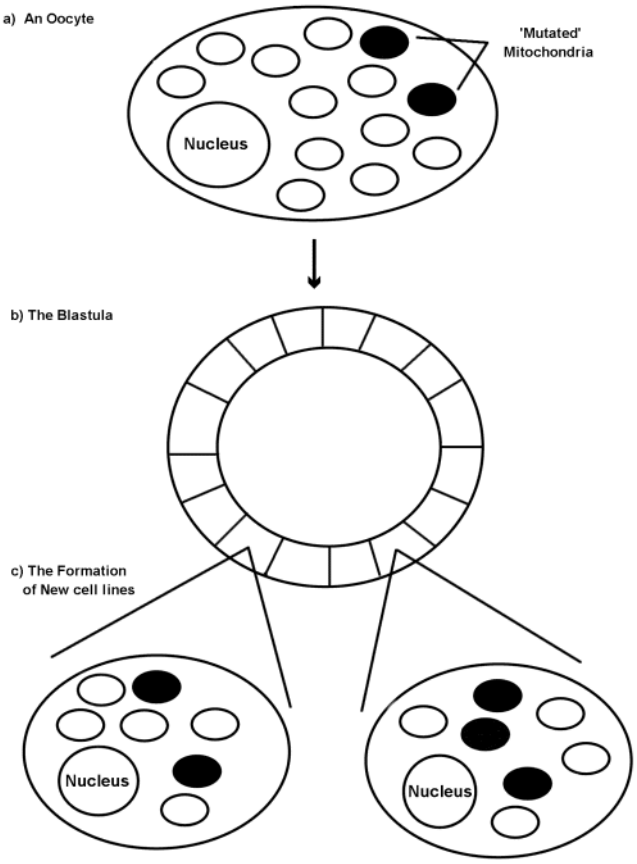

Figure 3. New cell lines containing different proportions of mutated mitochondria.

Figure 3 illustrates these stages in embryonic growth and shows how two cell lines may vary in the number of *mutated mitochondria* that they contain in their cells.

Heteroplasmy and Homoplasmy

The above discussion has described how it is possible for the cells in a particular cell line of the body to acquire a new mutation; and how if the cell line is the *Germ Cell Line* that produces the egg cells, the *oocytes*, for the next generation it is possible for the proportion of mtDNA molecules with the new mutation to be higher in the next, and subsequent generations.

At present it is not feasible to test cells from the *Germ Cell Line* for Genealogical purposes, so it is usual to test *buccal* cells which are very easy to collect using a mouth swab. A *buccal* cell sample is generally assumed to represent the cells of the *Germ Cell Line* - but as the above discussion has pointed out, there may be differences in the proportion of mutated mtDNA molecules in different cell lines.

The mtDNA molecules obtained in a cell sample can theoretically be considered as having *heteroplasmy* at very low levels, but it is really only possible when sequencing for Genealogical purposes for *heteroplasmy* to be detected when a new mutation is present in maybe a tenth of the mtDNA molecules, i.e. at a level of, maybe, 10%. This also means that *homoplasmy* is assumed on testing when the new mutation might actually only be present in about 90% of mtDNA molecules.

Just how a new mutation present at a low level of *heteroplasmy* propagates to become a *homoplasmic* mutation is not properly understood. But it is likely that in many cases where a new mutation is being propagated successfully that *heteroplasmy* levels rise slowly over a number of generations. This, however, does not mean that the level of *heteroplasmy* rises with every generation, and indeed the level can be expected to fall in some generations.

But whilst the gradual rise in levels of *heteroplasmy* over a number of generations is possibly the usual way in which a new mutation is propagated, this does not explain the observation that on occasions new mutations can be seen to appear with a high level of *heteroplasmy,* or even *homoplasmy,* in an individual, but the mutation is not found in the individual's mother.

However, one explanation of this phenomenon may be to consider that at a particular stage in the development of *Germ Cells* in an embryo there is such rapid growth that each cell contains very few mitochondria. And, if one of these mitochondria should contain a new mutation it is possible that the mutation may appear to propagate suddenly over a single generation. But, as discussed above, this does not mean the *initial mutational event* is recent, as the first phase of propagation leading to the presence of at least a single mutated mitochondria will have needed to have occurred beforehand.

The recent papers by Shoubridge and Wai (Shoubridge, 2007, Wai, 2008) discuss this phenomenon in some detail and describe it as being a *genetic bottleneck*.

The Computer Model accompanying this paper attempts to illustrate how a new mutation might propagate successfully from a very low level of *heteroplasmy* to *homoplasmy* - or, as happens in most cases, fails in its attempt to propagate. The model allows for the number of mitochondria per cell to be reduced to very low levels, thereby simulating the effect of a *genetic bottleneck*.

Three Sets of Example Inputs and Outputs

Set a --- With the default values.
Result - Heteroplasmy is achieved on 2 runs, but Homoplasmy is not achieved.

Initial Level of Heteroplasmy = 0.5%
Replications per Generation = 5
Level of Heteroplasmy over 20 generations
1. 0% ....
2. 0.5% 0% ....
3. 0% ....
4. 2.5% 8.5% 5% 2.5% 3.5% 4.5% 6.5% 8% 10.5% 13% 15% 11% 5.5% 3.5% 2.5% 5% 3.5% 5% 4% 4%
5. 0.5% 0% ....
6. 0% ....
7. 0% ....
8. 3.5% 1.5% 4.5% 2% 1% 1.5% 2.5% 1.5% 2% 3% 4% 8% 9% 7.5% 7.5% 11.5% 5.5% 14.5% 8.5% 6.5%
9. 0% ....
10. 0% ....
End of runs

set b --- Observing 1000 Runs (and suppressing the printing of unsuccessful runs).
Result - Homoplasmy achieved on 2 occasions.

Initial Level of Heteroplasmy = 0.5%
Replications per Generation = 15
Level of Heteroplasmy over 20 generations
462. 11.5% 20.5% 24.5% 31% 41.5% 61.5% 83.5% 100% HOMOPLASMY ACHIEVED
747. 1% 4.5% 4.5% 5% 8% 9.5% 10.5% 20% 36.5% 36.5% 56.5% 95.5% 100% HOMOPLASMY ACHIEVED
End of runs

set c --- the effect of reducing the number of mitochondria to a very low level.
Result - Homplasmy can be observed within a few generations on several of the Runs.

Initial Level of Heteroplasmy = 10%
Replications per Generation = 10
Level of Heteroplasmy over 20 generations
9. 50% 100% HOMOPLASMY ACHIEVED
17. 70% 90% 100% HOMOPLASMY ACHIEVED
48. 30% 60% 100% HOMOPLASMY ACHIEVED
50. 60% 40% 70% 100% HOMOPLASMY ACHIEVED
60. 20% 40% 100% HOMOPLASMY ACHIEVED
66. 50% 100% HOMOPLASMY ACHIEVED
72. 50% 80% 100% HOMOPLASMY ACHIEVED
93. 100% HOMOPLASMY ACHIEVED
98. 20% 90% 100% HOMOPLASMY ACHIEVED
End of runs

Phase B: Modelling the steps from Mitochondrial Capture to the appearance of Heteroplasmy and Homoplasmy

This part of the model shows how the level of heteroplasmy might vary from one generation to the next; and whether homoplasmy is achieved.

The starting condition can be considered as representing the presence of a single mitochondrion with a new mutation in all of its mtDNA molecules; and running the model shows over a series of **Runs** whether the states of heteroplasmy and homoplasmy are achieved.

In the computer model the user is asked to determine:

- the number of **Runs** from the range 10 - 1000,

- the **Mitochondrial Number** from the range 10 - 800. In effect this number fixes the initial level of heteroplasmy. Choosing '200' mitochondria will give an initial level of 1/200, or 0.5%.

- the number of **Replications** from the range 5 - 20. This number can be considered to represent the number of cell divisions at the 'blastula' stage. The higher the number selected the more randomisation will occur between generations.

- the number of **Generations** form the range 10 - 300. This number alters the number of generations that are to be followed.

- there is also the option of suppressing the display of the **Full Results,** which can be useful when the output from the model is large.

The model is run by choosing **B:START**. Results for three sets of example inputs are shown above.

Conclusions

The above discussion and accompanying Computer Model offer some ideas as to how mitochondrial mutations might arise and be propagated.

First, a mutation appears as an *initial mutational event*, before becoming the dominant form in the mitochondrial DNA molecules of a nucleoid, then a mitochondrion, before appearing in an individual as a *heteroplasmic* or *homoplasmic* mutation.

But if a particular mutation cannot be detected on sequencing as being *heteroplasmic*, does this mean the mutation is totally absent, or is it present in some of the mtDNA molecules of an individual? The ideas discussed above do suggest that the mitochondria of an individual may contain many new mutations, each of which has a very small chance of propagating over future generations to be a *heteroplasmic* or *homoplasmic* mutation.

However, other problems still remain to be solved:

For example, there are many examples in the phylogenetic tree where a new mutation appears in branches near to each other, but is absent in many other branches. The usual explanation of this phenomenon has in the past involved the concepts of *parallel mutations* and *back mutations.*

A *parallel mutation* being said to occur when the same mutation has appeared in different places in the tree, i.e. the mutation is said to have emerged by coincidence in different places. This phenomenon is generally considered simply to be the result of the *randomness of nature* as there is only a limited number of possible mutations that can occur and so it is likely the same mutation will occur by chance in different places in the phylogenetic tree.

And, a *back mutation* is said to occur when a new mutation is found higher up the tree, but is absent in lower branches where it would be expected. This phenomenon is difficult to explain as it would appear to depend on a *mutation* having a *memory* of its previous state--something that is unlikely to occur.

However, this paper describes how a new mutation must in the first instance appear as an *initial mutational event* and then may propagate to be a *heteroplasmic* or *homoplasmic* mutation. And, whereas the ideas described here do not explain fully the finding of *back mutations* in the phylogenetic tree, they do suggest a mechanism by which a new mutation may not appear in all the descendants of an individual.

There is clearly still a great deal to be learnt about the behaviour of mitochondrial mutations and better models will be produced, but the present paper together with its Computer Model appears able to explain many of the observed phenomena.

Supplementary Information

A listing of the code for the computer model is available both in text form and as a web page at::

http://www.jogg.info/51/nucloid_model.txt

http://jogg/info/51/nucleoid_model.html

References

Anderson S, Bankier AT, Barrell BG, de Bruijn MHL, Coulson AR, Drouin J, Eperon IC, Nierlich DP, Roe BA, Sanger F, Schreier PH, Smith AJH, Staden R, Young IG (1981) Sequence and organization of the human mitochondrial genome. *Nature*, 290:457-465.

Andrews RM, Hubacka I, Chinnery PF, Lightowlers RN, Turnbull DM, Howell N (1999) Reanalysis and revision of the Cambridge reference sequence for human mitochondrial DNA. *Nat Genet*, 23:147.

Cann RL, Stoneking M, Wilson AC (1987) Mitochondrial DNA and human evolution. *Nature*, 325:31-36.

Doublet V, Souty-Grosset C, Bouchon D, Cordaux R, Marcadé I. A thirty million year-old inherited heteroplasmy (2008) PLoS ONE. 2008 Aug 13;3(8)

Gilkerson RW, Schon EA, Hernandez E, Davidson MM (2008) Mitochondrial nucleoids maintain genetic autonomy but allow for functional complementation. *J Cell Biol*, 181:1117-28.

Kivisild T, Shen P, Wall DP, Do B, Sung R, Davis K, Passarino G, Underhill PA, Scharfe C, Torroni A, Scozzari R, Modiano D, Coppa A, de Knijff P, Feldman M, Cavalli-Sforza LL, Oefner PJ (2006) The role of selection in the evolution of human mitochondrial genomes. *Genetics*, 172):373-87.

Shoubridge EA, Wai T (2007) Mitochondrial DNA and the mammalian oocyte. *Curr Top Dev Biol*, 77:87-111.

Wai T, Teoli D, Shoubridge EA (2008) The mitochondrial DNA genetic bottleneck results from replication of a subpopulation of genomes. *Nat Genet*, 40:1484-8.