Mathematics and Modeling in Cryonics: Some Historical Highlights

By R. Michael Perry and Aschwin de Wolf

Introduction

Robert Ettinger, in *The Prospect of Immortality*, his 1964 classic that largely started the cryonics movement [34], speaks of "the solid gold computer" which is to herald a "second industrial revolution." This in turn will rest "on the replacement of human brains by machines." Marvin Minsky, one of the founders of artificial intelligence (and a reported cryonics patient [23,37]) is among those quoted in support: "I believe … that we are on the threshold of an era that will be strongly influenced, and quite possibly dominated, by intelligent problem-solving machines."

Computers are, of course, intimately associated with mathematics, for what is a computer but a machine for the automated doing of mathematics and mathematical modeling in their many, varied forms, in application to problems of interest? This report will not directly concern the use of computers in cryonics, but the underlying mathematical initiatives that have occurred over the years. First came theoretical work, then practical implementation, again using computers. Cryonics has been, of course, a small movement sponsored by private interests. It is not the place where you expect to find armies of scientists and mathematicians, armed with the latest computing technology, applying algorithms on a massive scale and also developing new ones at breakneck pace. That said, the extent of mathematical investigation and application in cryonics is rather surprising, a bit more than you'd expect from the labeling of "pseudoscience" that those outside the field often attach to the practice [16].

Here we look at several efforts over the years to address cryonics problems mathematically, using layman's (mostly non-mathematical) language. We start with pioneering work of Art Quaife with his remarkable modeling in the early 1970s of perfusion systems, with a follow-up in the 1980s. We summarize Hugh Hixon's work using the Arrhenius equation to gauge the adequacy of cryopreservation temperatures. An effort of the author at modeling perfusion systems is then recounted, with use on an Alcor case. Next is a different problem: estimating the total post-pronouncement ischemic exposure of cryonics patients, until they are cooled sufficiently that further ischemic challenge is effectively halted. Some preliminary work by the author is summarized, followed with a more thorough if unpublished treatment by Steve Harris. Next is the work of Greg Fahy and others at 21st Century Medicine to predict cryoprotectant toxicity and thus shorten the labor of finding better cryoprotectants. A different problem then covered is to estimate the expected frequency of heavy caseloads, given that cases essentially occur randomly. Next is the work of Robert Freitas to address the problem of long-term organizational stability through econometric modeling. Our coverage concludes with some work by Aschwin de Wolf and others, which focuses on ischemic damage in neural tissue, including an algorithm implemented by Michael Maire to characterize ischemic damage based on machine learning.

Art Quaife's Early Perfusion Modeling Paper

Fred and Linda chamberlain, who founded Alcor in 1972, also established a sister organization, Manrise Corporation, which would deal with technical matters relating to cryonics procedures. Its journal, *Manrise Technical Review* (MTR), ran through seven issues 1971-73, crammed indeed with technical material, both theoretical and practical, on procedures for cryopreservation, as well as a detailed report on an actual case. Coming so early in cryonics history (within 10 years of the first freezings [31]) the journal stands as a remarkable tribute to the tenacity and seriousness with which the idea of cryonics was being pursued by some, as well as the high level of competence of some of those pursuing. Art Quaife's paper, "Mathematical Models of Perfusion Processes" appeared in the April-June 1972 MTR [4] and runs for dozens of finely printed, equation-heavy pages. It is a tour de force of mathematical insight into perfusion and the closely related heat-exchange problem that should still be of interest to the mathematically inclined. The abstract notes the pioneering nature of the work, then briefly summarizes the content. "Solid state hypothermia" (SSH) refers to the state of being at cryogenic temperature (usually liquid nitrogen temperature):

> Very little empirical data is available concerning the actual efficacy of alternate procedures for freezing and thawing whole mammals to[/from] the temperature of liquid nitrogen. In order to currently formulate recommended procedures for the induction of solid state hypothermia in humans, it is thus necessary to rely partly upon extrapolation from the results of experiments treating the freezing of cells and organs,

> and to an even greater extent upon theoretical analysis of the effects of alternate procedures. In this paper, we formulate mathematical models of many of the processes that take place during the induction of SSH in humans. Primary focus of the treatment is to analyze the rate of removal of heat from the body, and the rate of buildup of cryoprotectant in the body cells during perfusion. Specific formulas are also developed giving the quantity of cryoprotectant and the length of time required to accomplish perfusion, the cooling profile to be followed, and many others. The conclusions of the theory are used to formulate a specific recommendation as to a best current perfusion procedure. In many cases, the unknown values of the body parameters appearing in the equations require that approximating assumptions be made in order to achieve numerical results. It is expected that these results will be refined considerably when further information becomes available concerning the values of these parameters.

Some additional detail is added in the short Introduction:

> As currently conceived, the induction of solid state hypothermia (SSH) in the human body takes place in three phases. During Phase I, the blood is replaced by a suitable balanced salt solution and the body temperature is lowered to 0°C as rapidly as possible. In Phase II, a concentration of cryoprotectant [assumed in the paper to be dimethyl sulfoxide (DMSO)] is built up in the body water, while continuing to lower the body temperature in such a way as to remain in liquid state. Phase III begins when the desired terminal concentration is achieved; the body is then lowered to liquid nitrogen temperature where it is stored in solid state. During the first two phases, systems will be required to recirculate chilled perfusate through the body. This paper treats the mathematics of such systems.

A preliminary report on this work appeared in *The Outlook* (April 1972), before the publication in MTR [6]:

> …Mr. Quaife has gone well beyond any known previous work in applying detailed analysis to the problems involved, including the major questions of cooling rates and control of concentration of perfusate and its ingredients, and the related questions of temperature gradients, concentration gradients, minimizing of times and costs, etc.
>
> The degree of sophistication, both biological and mathematical, is considerable. Differences between organs and tissues are considered, and formulation of some problems leads to differential equations to be solved by the method of the Laplace transform. Nevertheless, it is hoped that as many as possible will read the paper, get what they can out of it, and feed back any useful comments. A revised draft will be published in an early issue of *Manrise Technical Review*.

The "revised draft" is, in fact, the paper under consideration here. The "three phases" Art Quaife refers to in the cryopreservation process ("induction of solid state hypothermia") are now called body washout (Phase I), cryoprotective perfusion (Phase II) and cryogenic cooldown (Phase III). Phases I and II involve perfusion of substances (solutions) into the body and Phase III takes the patient from approximately water ice temperature to the cryogenic range, typically the temperature of liquid nitrogen. The paper mainly concerns these first two phases, where substances are being perfused into the body and (initially) blood and body fluids are replaced. The perfusion makes use of the body's vascular system, otherwise used when the heart circulates blood. The performance of the vascular system is highly tissue-dependent, as Art explains:

> In normal circulation, there exists a wide variation in the percentage of the cardiac output delivered to different tissues. For example, on a per gram of tissue basis, the kidneys receive over 150 times as much blood as do the skeletal muscles. We can approximate the normal blood flow by dividing the body into two parts: (a) Strongly circulated tissue, consisting of the brain, heart muscle, kidneys, and organs of the hepatic-portal circulation (mainly the liver), and (b) Weakly circulated tissue, consisting of the remainder of the body.

The strongly circulated tissue makes up a much smaller volume and mass than the weakly circulated, yet the total circulation through the strongly circulated is several times greater. A cited reference estimates that only 7.3% of the body weight of an adult human male (63 kg) is strongly circulated, at 3.76 liters/min. (resting state of cardiac output, as are the other estimates here). The balance is weakly circulated, at 1.64 liters/min., for a total circulation of 5.40 liters/min. For the case of perfusion of a postmortem body for cryonics purposes a figure of 4 liters/min. total circulation is assumed, with the other values scaled proportionally (2.79 and 1.21 liters/min. respectively). In this way we obtain the "partitioned model" of the body comprising the two different flow rates in the different, nonoverlapping tissue components. Also considered is the "uniform model" in which the circulation flow rate is assumed to be the same throughout the body, and still the same overall (4 liters/min.).

As noted above, in Phase I perfusion, blood and body fluids are replaced with "base perfusate" which sets the stage for Phase II, cryoprotection. An additional effect in Phase I is to cool the body, typically from near normothermic temperature (37°C) to near the ice point (0°C). The chart below ("figure 4" in the paper) shows predicted effects (calculated temperature curves) if chilled perfusate at 0°C is circulated through tissue starting at a temperature of 37°. Shown are the three cases of (1) uniform

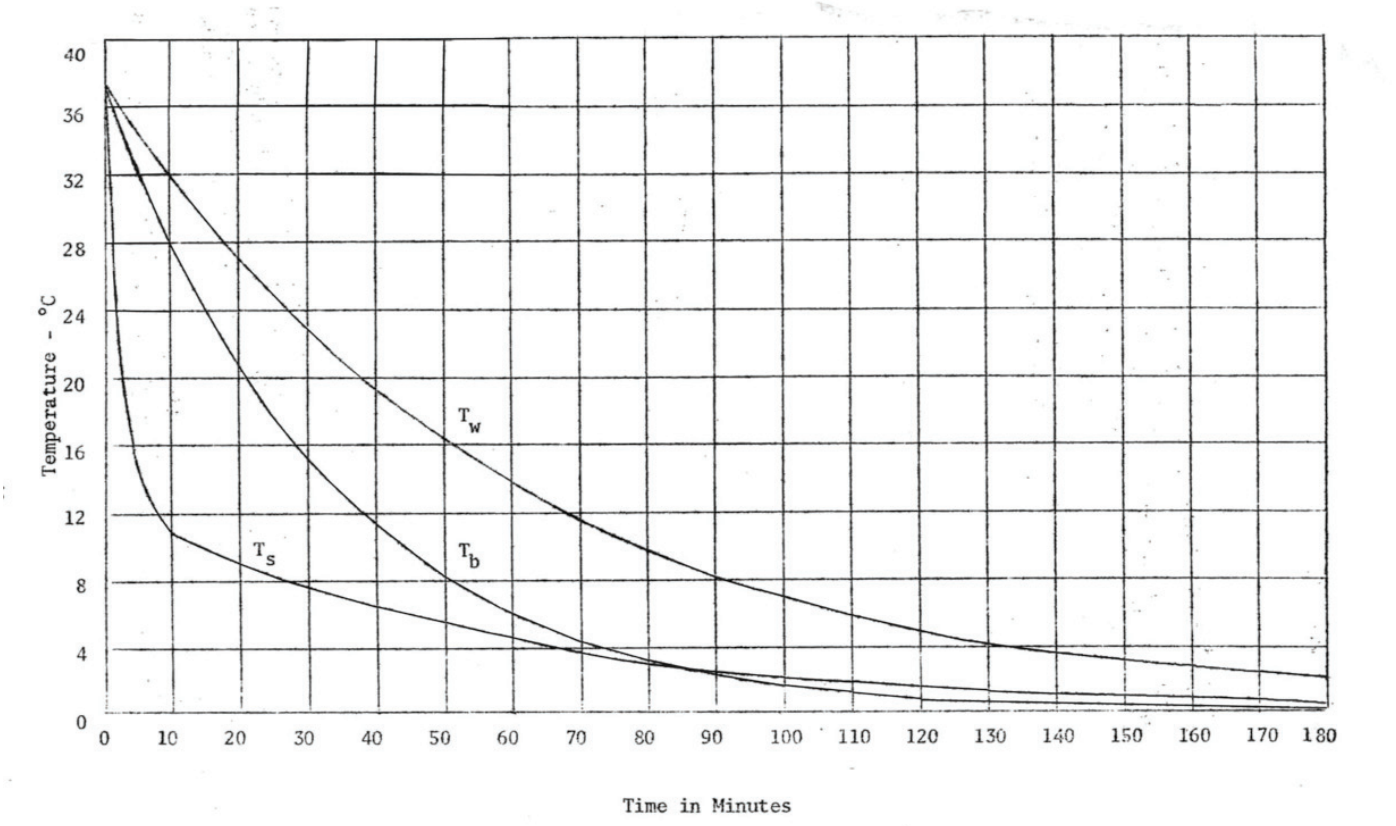

Predicted cooling curves for a 63 kg. man, with perfusion flow of 4 liters/minute and heat exchanger efficiency H = .4. T_s and T_w are from the partitioned body model, while T_b is from the uniform body approximation. Due to the limitations mentioned in the text, the actual rates of cooling of T_s and T_w probably lie closer to T_b than do the predicted curves.

body model (T_b), (2) strongly circulated portion (T_s), and (3) weakly circulated portion (T_w). The caption is taken from the paper, where it is noted that, due to various uncertainties, the true values may differ significantly from predictions.

Predicted cooling curves for a 63 kg. man, with perfusion flow of 4 liters/minute and heat exchanger efficiency H = .4. Ts and Tw are from the partitioned body model, while Tb is from the uniform body approximation. Due to the limitations mentioned in the text, the actual rates of cooling of Ts and Tw probably lie closer to Tb than do the predicted curves.

There is much more to the paper. Not covered here, for instance, are detailed recommendations of perfusion circuit designs and suggested flow rates during the different phases of the cryoprotection, and how to determine them. Some further, if still highly abbreviated, idea of the scope and depth of the work can be gathered from the titles of the nine main sections: (1) Limitations of the Analysis; (2) Macrocirculation of the Blood; (3) Microcirculation of the Blood; Diffusion and Heat Conduction; (4) Black Box Description of Heat Exchangers and of the Body as a Medium for DMSO Diffusion; (5) Cooling the Body; (6) Phase I Perfusion; (7) Cryoprotection; (8) Freezing Point of DMSO-Water Solutions; (9) Phase II Perfusion.

Cryoprotectant mixtures have changed since the date of this study, when DMSO (dimethyl sulfoxide) was the main ingredient [12,13,22], yet the basic approach and mathematical analysis remain valid and worthy of study.

Second Art Quaife Paper

Fast-forwarding a few years, a second Quaife paper appeared, "Heat Flow in the Cryonic Suspension of Humans: Survey of the General Theory" (*Cryonics*, September 1985) [6]. The problem of heat flow in solids has been extensively studied in physics, and Quaife offers some textbook equations that describe the general problem before proceeding to some simplified versions of the problem that approximate conditions in human cryopreservation. The abstract offers further details:

> Procedures used in the successful freezing and thawing of diverse human cells and tissues are known to be quite sensitive to the cooling and thawing rates employed. Thus it is important to control the temperature descent during cryonic suspension of the whole human body. The paper surveys the general theory of macroscopic heat flow as it occurs during the cryonic suspension of human patients. The basic equations that govern such

heat flow are presented, then converted to dimensionless terms, and their solutions given in geometries that approximate the human torso, head, and other regions of the body. The solutions are more widely applicable to the freezing of tissues and organs.

The article, then, generalizes the results of the earlier paper, then specializes the treatment to approximate cases of interest in cryonics. It is divided into thirteen main sections as follows: (1) Nomenclature; (2) Equations Governing Heat Conduction; (3) Dimensionless Variables; (4) Initial and Boundary Conditions; (5) Dimensionless Heat Flow Equations; (6) Global Reformulation of Equations; (7) Useful Mathematical Functions; (8) Basic Solution: Heat Flow in One Dimension; (9) Heat flow in a Semi-Infinite Solid; (10) Heat Flow from a Highly Insulated Solid; (11) Heat Flow from a Sphere; (12) Heat Flow from a Cylinder; (13) Analogy between Heat Flow and Diffusion.

The brief Introduction, quoted here, informs the reader about cryonics then notes how the present work extends and generalizes work completed earlier (above). Calculations are promised for later articles, though not published as far as we are aware:

Cryonic suspension is the freezing procedure by which human patients are preserved, after pronouncement of legal "death," in hopes of eventual restoration to life and health. The procedure attempts to preserve the basic information structures that determine the individual's identity. These include the memories and personality as encoded in the macromolecules and neuronal weave of the brain, and the genetic information stored in DNA.

The author has previously formulated a mathematical model of the heat flow and the diffusion of cryoprotectant that occurs during the first phase of this procedure, in which chilled blood substitutes and cryoprotective solutions are perfused through the vascular system. The present paper treats the general theory of heat flow, particularly at sub-zero temperatures after perfusion has ceased and the body has solidified.

The author has written a computer program that calculates most of the solutions given below, and in subsequent articles intends to present tables and graphs comparing theoretical projections with experimental data. Other problems for subsequent analysis include

HEAT FLOW IN CRYONIC SUSPENSION

HEAT FLOW		⟷	DIFFUSION	
Quantity	*Unit*		*Quantity*	*Unit*
ρcT	J/m^3		C	kg/m^3
α	m^2/s		D	m^2/s
q	$J/(m^2\,s)$		j	$kg/(m^2\,s)$
F	J/m^2		M	kg/m^2
Q	J		m	kg
$h/\rho c$	m/s		α	m/s
Equation	*Name*		*Equation*	*Name*
$q = -k\nabla T$	Fourier's law		$j = -D\nabla C$	Fick's first law
$\nabla^2 T = \frac{1}{\alpha}\frac{\partial T}{\partial t}$	Fourier's equation		$\nabla^2 C = \frac{1}{D}\frac{\partial C}{\partial t}$	Fick's second law
$k\frac{\partial T}{\partial t} = -h(T - T_\infty)$	boundary condition		$D\frac{\partial C}{\partial t} = -\alpha(C - C_\infty)$	boundary condition

Table of correspondence for heat flow versus diffusion

> change of phase, and the thermal stresses from temperature gradients within the frozen tissue.

Next, after variables are defined and terminology is established, basic equations are given: heat flow in a solid, average temperature over a volume in space representing an object such as a cryonics patient, average temperature over the surface of the volume in question, and so on. The general equations are, of course, "textbook physics" and can be found in standard references but are then specialized to address matters of interest in cryonics. Of particular interest is the treatment of heat flow in a sphere and a right circular cylinder, which approximate a human head and torso, respectively. In both cases solutions are greatly simplified over the general case, both from the simple geometry and by assuming uniform thermal conductivity which doesn't vary with temperature. (The cylinder is also assumed to be infinite in both directions, admittedly a bit of a "stretch" for a real patient, but presumably not very different inside except near endpoints.)

The work is, in fact preliminary. No calculations are offered, only formulas. That said, it is a very good start, and going further and implementing an actual modeling of a patient would be feasible but daunting and has not been attempted yet or at any rate nothing is published about such an attempt, as far as we are aware.

The paper primarily deals with heat flow, which is certainly important in cryogenic cooling, but leaves aside another problem, that of inducing cryoprotectant in desired concentration before deep cooling begins. Here, however, physics is emphatically on our side, as Quaife noted in his earlier study and reminds us here: Diffusion of a substance in a fluid medium (where mass is being transferred rather than energy in the form of heat) is mathematically equivalent to heat flow in a solid. As a parting shot Quaife in his second paper provides a table of correspondence between the two.

How Cold Is Cold Enough?

Cryonics depends on the idea that, if you store a biological sample at a low, subfreezing temperature it will be essentially unchanged for (at least) a few centuries. But how can you reliably estimate the rates of change in materials at different temperatures? And how cold do you have to be that the rates you are interested in are low enough that you can stop worrying? Biochemist Hugh Hixon, a long-time Alcor staff member still employed there today, set out to answer this question in the mid-1980s ("How Cold is Cold Enough?") [21]. His approach was to use the well-known Arrhenius equation to estimate and compare chemical reaction rates at different temperatures.

Though a good start, this approach encounters a major difficulty: in a tissue sample there are many chemical reactions going on, and their rates, which might be well-approximated individually by the Arrhenius equation, vary widely. So, for instance, a certain reaction might slow down by a factor of two when you drop the temperature a certain amount, while another might slow down by a factor of five. Hixon addresses the difficulty by selecting one important substance, catalase, and basing all conclusions on its reaction rate as a function of temperature:

> I am going to be pessimistic, and choose the fastest known biological reaction, catalase. I'm not going to get into detail, but the function of the enzyme catalase is protective. Some of the chemical reactions that your body must use have extraordinarily poisonous by-products, and the function of catalase is to destroy one of the worst of them. The value for its E [important in the Arrhenius equation] is 7,000 calories per mole-degree Kelvin. It is sufficiently fast that when it is studied, the work is often done at about dry ice temperature. My friend Mike Darwin remarks that he once did this in a crude fashion and that even at dry ice temperature things get rather busy. Another reason to use it is that it's one of the few I happen to have. E's are not normally tabulated.

With this choice, it is then possible to compare reaction rates as the temperature is lowered from body temperature (37°C) on down. And the results, overall, are reassuring, at least for the usual cryogenic storage. What takes one second to happen at body temperature would take about 25 million years at the temperature of liquid nitrogen! For temperatures warmer than liquid nitrogen the rates will be faster, but the tentative conclusion is that -115°C would slow things down enough that 100 years of storage would be equivalent to 12 hours at body temperature, at least barely acceptable. (The point is made too that below about -135°C translational molecular motion is inhibited so safe storage of almost indefinite length should be possible, irrespective of the Arrhenius equation.)

Perfusion Modeling

The next study was done by Perry ("Mathematical Analysis of Recirculating Perfusion Systems, with Application to Cryonic Suspension," *Cryonics*, October 1988) [24]. Less comprehensive and general than the Quaife articles, it deals mainly with cryoprotective perfusion (Quaife's Phase II), but also reports some use with an actual case, where predicted and measured values of cryoprotectant concentrations are compared. Here the patient's temperature has been lowered to near the ice point (0°C) and the base perfusate introduced in Phase I is replaced with cryoprotective agent (CPA) to protect the tissues during cooldown to cryogenic temperatures (Phase III). The conclusion of the article summarizes the main rationale and results of the study:

> Cryonic perfusion, undertaken to protect tissues from damage during the freezing process, is nonetheless

not a completely benign process. Cells and tissues can suffer damage during perfusion from (1) toxicity of cryoprotective agent(s), (2) osmotic stress, and (3) stress resulting from mechanical forces under excessive fluid pressure. For this reason the perfusion process must be carefully controlled so that CPA is introduced rapidly but without extreme concentration gradients or excessive pressure. Perfusion circuits currently in use by Alcor allow careful control of CPA buildup through control of the flow rates into and out of the patient and other fluid reservoirs involved in the perfusion. How to control the fluid flow rates to achieve an effective perfusion is a complex problem, but one that can be addressed through mathematical modeling of perfusion circuitry. It thus becomes feasible to predict with reasonable accuracy the rate of increase in CPA concentration for cryonic perfusion circuits now in use. The method shows promise in elucidating what is happening to a patient during perfusion, in comparing different perfusion protocols [based on] quantities related to cell and tissue stress, and in selecting protocols to achieve optimal perfusion under given models of stress.

In the paper, a perfusion system consisting of n reservoirs is modeled with arbitrary, pairwise interconnections and flow rates that are constant with time. A binary solvent-solute mixture is circulated through the reservoirs, with mixing of incoming fluids in each reservoir assumed to be instantaneous. We then wish to know the volume-for-volume (v/v) concentration of solute in each reservoir as a function of time. This in turn depends on a linear, ordinary differential equation involving flow rates into and out of the reservoir. The equation is solved numerically, using a Taylor's series approximation, and is then generalized to the case in which the flow rates vary with time.

Application of the technique to the problem of CPA perfusion during a cryopreservation is then considered. For this case we have 4 reservoirs (n = 4): the patient, a concentrate reservoir, a recirculating reservoir, and a discard. (See illustration.)

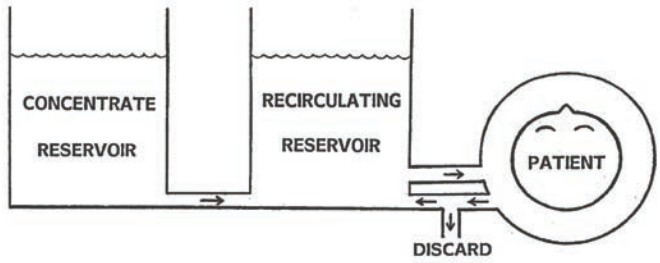

Perfusion circuit as modeled in the paper has 4 reservoirs including the patient and discard.

The direction of flow of perfusate for the different reservoirs is as follows: (1) concentrate: 1-way, to recirculating; (2) recirculating: 2-way, to/from patient; (3) patient: 2-way, to both recirculating and discard, from recirculating; (4) discard: 1-way, from patient. The two most important flow rates are (1) flow rate into patient from recirculating reservoir (f_{IN}) and (2) flow rate from patient to discard (f_D). The difference f_{IN}- f_D is the rate of flow from the patient back to the recirculating reservoir, assuming no change in volume in either the recirculating reservoir or the patient, so that f_D is also the flow rate from the concentrate to the recirculating reservoir.

In the course of perfusion, the concentration of solute in the recirculating reservoir and then the patient increases. In an actual case the perfusate is circulated in the patient's vascular system, divided into arterial (ingoing) and venous (outgoing) components. The arterial-venous or "a-v" difference in the two concentrations is a measure of osmotic stress: a greater difference signifies a greater stress, which we would like to reduce as far as possible, consistent with avoiding excessive flow rates that might cause edema. Reduction occurs by recirculating the perfusate between the patient and the recirculating reservoir as shown in the illustration. In absence of recirculation (perfusate discarded from patient, no inflow back to recirculating reservoir) we would obtain a "one pass" perfusion circuit with correspondingly larger a-v differences.

Putting it all to the test

A computer program was written to predict the concentration of CPA in the patient and elsewhere as a function of time, assuming constant flow rates into and out of the patient and the reservoirs involved in perfusion. (In practice the only other case of interest that requires calculation is the recirculating reservoir.) The program was used to simulate perfusions, showing how recirculation improves the quality of perfusions by reducing a-v differences.

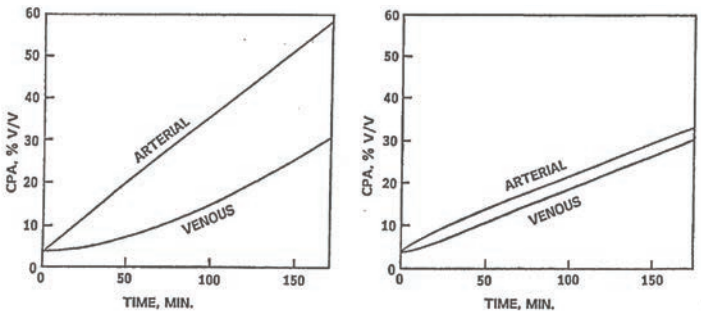

Simulated perfusions. Left is single-pass (no recirculation) and shows much greater a-v differences than right which uses recirculation under otherwise similar conditions.

The program was then used during a whole-body cryopreservation at Alcor (James Binkowski, starting 8 May 1988) [32,33]. Rough estimates were obtained, in advance, of time and perfusate requirements under different assumed flow rates, and actual flow rates were adjusted accordingly to reduce

anticipated osmotic stress. A more careful study was done afterward to try to reconstruct the course of CPA perfusion. Comparison of predicted and observed CPA concentrations shed light on certain physical changes, such as reduction of patient circulating volume, believed to have been caused by vascular occlusion secondary to ischemic clotting. Other programs were then written that allowed for variable flow rates, to obtain optimal perfusion profiles under certain models of cell stress. Calculations from one of these programs are shown below.

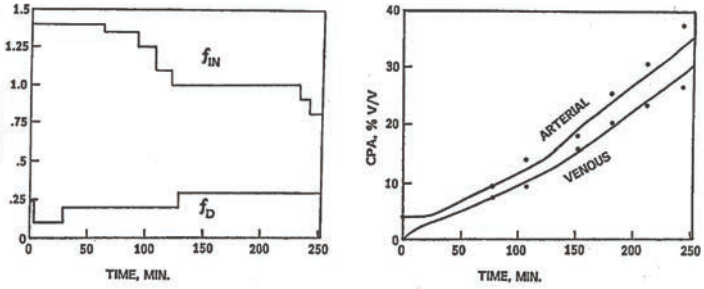

Left: flow rates for actual perfusion.
Right: observed arterial-venous CPA concentration (dots), and values calculated from model (solid curves).

A section of the paper on neuropreservation showed how this procedure would lead to a reduction in osmotic stresses over the whole-body case, under starting assumptions that seemed likely at the time. It was hoped that further programming efforts might lead to improvements including even a computer-controlled system for cryoprotection. This did not materialize, however, and use of the program was discontinued when it appeared that hands-on experience was an adequate guide for directing the course of cryoprotection. (At the time the study was carried out the CPA in use was glycerol. This use has been superseded [13], but, just as in the Quaife studies based around DMSO, the mathematics could apply to other choices.)

Another complicating factor is the effect of ischemia on flow rates and tissue equilibration. As the duration of ischemia increases, perfusion of the brain is progressively compromised due to the so called "no-reflow" phenomenon. For a given pressure (say 100 mmHg), flow rates will slow down significantly in ischemic patients, increasing perfusion times. Another adverse effect of ischemia is the development of swelling (edema) and intracranial pressure. Under such conditions a very smooth increase in cryoprotectant may not be desirable because it leads to significant water accumulation in the cells and tissues before exposure to the highest concentrations. It is also believed that a steeper increase of the cryoprotectant can be beneficial because the higher osmotic difference can recruit (edematous) water back into the circulatory system. For example, in an ischemic patient cryoprotection would not start at 0% CPA but 5% CPA. The understanding of various perfusion protocols on the ischemic patient is still rudimentary and no formal models are available yet.

Estimating Ischemic Exposure

Next we consider a problem in cryonics that is not related to perfusion or heat flow per se yet is certainly a vital concern: ischemic exposure of the patient prior to deep cooling. Historically, this problem was considered by Perry in an article in *Cryonics* (2nd Q 1996) [26] and later treated at much greater length in an unpublished (and incomplete) study by Steven B. Harris, MD [39]. Still later, Perry and de Wolf considered the problem of what cooling rate would be needed to escape ischemic injury, based mainly on Perry's earlier results [8].

What follows is adapted from the *Cryonics* article with an extension to summarize the Perry and de Wolf and Harris work. The Harris term E-HIT has been substituted for the author's original MIX ("measure of ischemic exposure") and his normalization is used. (E-HIT of 1 means 1 hour of exposure at body temperature of 37°C.)

An ever-present unknown in cryonics is the quality of a cryopreservation. Until more is known, in fact, we will have no good assessment of preservation quality in terms of what we'd really like to know: how well memory and other identity-critical elements are preserved in the cryopreserved remains. Meanwhile we are interested in whatever reasonable indicators of preservation quality it may be feasible to compute, though acknowledging these are imperfect.

One such possible indicator would be an "equivalent homeothermic ischemic time" (E-HIT) intended to assess the amount of high-temperature exposure the patient experiences, mainly (not exclusively) in the early stages of cryopreservation before the temperature falls below the ice point (0°C). Basically, the E-HIT would tally up how long, in hours say, the patient has been at a given temperature, with a heavier weighting used for higher temperatures, since more damage is occurring at these temperatures. According to a rule of thumb in wide use, each decrease of 10°C is supposed to cut in half the amount of damaging activity. At least this is considered roughly accurate – though it must not be pressed very far. Here this "Q-10" rule is adopted with the understanding that it is only a starting point.

So, if we normalize our measure so that an E-HIT of 1 corresponds to one hour at body temperature (37°C), then, one hour at 27°C would be an E-HIT of only 0.5, one hour at 17° only 0.25, and so on. More generally, for a temperature T in °C the E-HIT for one hour exposure is one-half raised to the power of one tenth the difference between body temperature and T, measured in degrees C, or in mathematical notation, $(1/2)^{(37-T)/10}$, in units of hours. We also assume that the damage scales linearly with the time of exposure; two hours at a given temperature would produce twice the amount of E-HIT as one hour at the same temperature, and so on.

So far we have considered the E-HIT if the patient is maintained at a constant temperature for an interval of time, but with a variable temperature we divide the time into small subintervals of approximately constant temperature and add up the contributions of each subinterval to obtain the total E-HIT. Mathematically, temperature is now a function of time *t*, that is *T*(*t*), and we must integrate the corresponding (instantaneous) E-HIT function, $(1/2)^{(37-T(t))/10}$, between two time limits defining the start and end of the time interval where cooling occurs, to obtain the total E-HIT.

Although this can get complicated, it turns out that the E-HIT takes a particularly simple form for the case of a constant cooling rate: the E-HIT for this case is just inversely proportional to the cooling rate. The proportionality constant will depend on particulars such as the time interval, the value assumed for body temperature, exactly how much reduction in damage rate occurs per unit drop in temperature (simple variants of the Q-10 rule, more or less than a 50% reduction per 10° drop, are possible), and so on. The chart below shows the cooling rate needed to achieve a given value of E-HIT, assuming the Q-10 rule, with cooling proceeding from body temperature to 3°C, approximately where cryoprotective perfusion typically occurs. As expected, reducing the E-HIT requires a faster cooling rate, while a slow cooling rate will produce a large E-HIT because a long time will be needed to achieve the targeted drop in temperature.

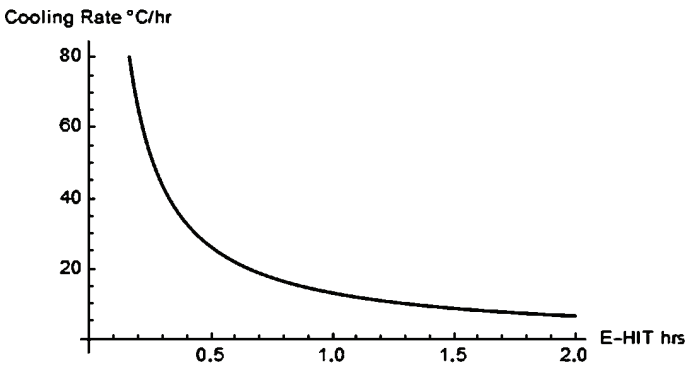

Cooling rate needed to achieve a given E-HIT, for constant-rate cooling from 37°C to 3°C, assuming the Q-10 rule.

It is instructive to consider something like a typical cryopreservation under presumed good conditions to see what the total E-HIT would be. Our model cooldown will be in three stages: (1) initial cooldown from body temperature to 3° at a "typical" rate of 15°/hour; (2) pause at 3° for 4 hours (= 240 minutes, close to the case of the previous section) for cryoprotective perfusion; (3) resumption of cooldown, again at 15°/hour, from 3° down to 123.3°, the glass transition temperature of cryoprotected tissue, where we assume that further deleterious change is halted. With these assumptions the total E-HIT will be 1.34 hours, with the breakdown as follows: stage 1: 0.87; stage 2: 0.38; stage 3: 0.09. (Actually, not much is known about the real contribution to E-Hit of cooling below 0°C, but it is expected to not be much larger, and maybe much smaller, than what follows from the Q-10 rule, and small in any case for reasonable cooling rates.) We see then that most of the E-HIT comes from the initial cooldown to the start of cryoprotective perfusion. A lesser amount, order of half as much, happens during the perfusion, while a much smaller amount still is from the further cooldown to the cryogenic range.

It should be emphasized that the E-HIT as we have computed it rests on the assumption that the patient is not ischemic prior to the start of cooling, that the cooling rate is uniform as indicated, and that no assistance such as metabolic support is used to lessen the E-HIT. Another point worth making is that, of course, one would like as small an E-HIT as possible, which raises the question of how small an E-HIT might have to be to be considered insignificant and not a matter of concern. Here, however, the answer is not particularly encouraging: An E-HIT of just five minutes (0.083 h) will begin to produce brain damage, as occurs in people who undergo a delayed resuscitation after cardiac arrest. Ideally, then, we'd like our total E-HIT not to stray outside the five-minute limit; unfortunately, this is not feasible and in fact E-HITs far in excess of this limit are the best we can achieve, as suggested by the example just considered. (More specifically, it would take a speedup factor of 16.1, or a cooldown rate of 4.0°/min., plus a similar reduction by a factor of 16.1 in the perfusion time at 3°C to just 14.9 min., to achieve an E-HIT of only 5 minutes! Basically, similar results were noted by Perry and de Wolf using somewhat different assumptions about the cooling/perfusion protocol, with infeasible cooling rates needed even skipping the pause at 3° for cryoprotective perfusion.) While this exercise has shown quite dramatically that an E-HIT of 5 minutes cannot be achieved through cooling alone, cryonics organizations employ other means to protect the brain during stabilization. For example, the results in an "ideal" cryonics case are more favorable because of the use of (mechanical) cardiopulmonary support (CPS) and the administration of cerebroprotective medications. Further work needs to be done to incorporate the effect of restoring circulation to the brain in these calculations. For example, can the E-HIT number be reduced by 50% if the patient is oxygenated during CPS and blood substitution?

We note here in passing some research directed toward reducing the initial whole-body cooldown time (body temperature around 3°C), in which mathematical modeling is used. Currently the best method is still "cardiopulmonary bypass," CPB, in which the blood and body fluids are circulated through a heat-exchanger and oxygenator to maintain metabolic support while cooling occurs. Alternatives that have been tested involve circulating cold fluid through the lungs ("liquid ventilation" or "total liquid ventilation," TLV) and a variant using gas as well as liquid ("gas/liquid ventilation," GLV). A study by Steve Harris, M.D. et al. using a canine model achieved cooling rates comparable

to those of CPB with a much less invasive, GLV procedure [20], though the technique has not yet found use in cryonics. However, the study is also notable for elaborate mathematical modeling based on Harris's study of the cooling process, which we now consider.

Harris in his unpublished work recognized that the assumption of constant-rate cooling is often unrealistic. Instead what is commonly encountered is that the cooling rate itself is a decreasing exponential. That is, one has a starting temperature T_0 and a target T_1, and after waiting time w the temperature one is at, call it T_i, is half-way between T_0 and T_1, and after that same additional time (a total of $2w$) it is half-way between T_i and T_1, and so on. True, at this rate the target temperature T_1 will never be reached but the approach can be quite close after only a few multiples of w, or T_1 might actually be lower than the desired target, which in turn could be reached quickly. The assumption of exponential cooling rates leads to a more complicated, but still computationally manageable, determination of E-HIT which better tracks what is really happening. Dr. Harris in his extensive study also offers insight into reducing E-HIT by such means as metabolic support. At present, though, the work on E-HIT is still preliminary, and more work is needed to produce a trusted indicator of brain injury or stress during the cryopreservation process.

The development of a quantitative outcome measure does not need to be confined to the induction of hypothermia. CT scans can be used to develop a score corresponding to the amount of ice formation observed in the brain. These scores can be combined to create a comprehensive outcome metric for a specific cryonics case. For example, if we limit ourselves to eliminating ischemia and ice formation, a "perfect" cryonics case would be one with an E-HIT of 0 (minutes of normothermic ischemia) and no ice formation in the brain (as inferred from CT scans).

Predicting Cryoprotectant Toxicity

Our next example is a simple one mathematically, but in other ways rather involved: to estimate the toxicity of a cryoprotectant mixture (CM) to be used in the cryoprotective perfusion step of the cooldown process (see above). *This step is important*. With today's protocols, such as the one now in use at Alcor, the tissue on further cooling to cryogenic temperature enters a glassy or "vitrified" state in which damaging ice crystal formation does not occur and damage overall, by appearance, is greatly reduced. Unfortunately, there is a tradeoff: known CMs also are toxic. Just warm the tissue up from a vitrified state and it either does not resume function, or if it does, its function is likely to be impaired except in the case of small tissue samples such as rat neonatal intestine, blood vessels, cartilage, corneas, or mouse ovaries [38,14]. Finding a CM that has minimum toxicity but still vitrifies at realistic cooling rates is thus a priority. How does one go about doing this? Experimenting with different possible CMs is a laborious process. Is there any way to shorten the labor?

Work of cryobiologists Gregory Fahy, Brian Wowk and others at 21st Century Medicine (21CM) uncovered a method of predicting the toxicity of cryoprotective solutions to a fair accuracy based on their molecular constituents. Results are detailed in a 2004 *Cryobiology* article [17], where the authors summarize their work as follows:

> The mechanisms of toxicity of vitrifiable solutions have not been elucidated. In part for this reason, it is not presently possible to predict the toxicity of either individual cryoprotective agents or mixtures thereof, and there is a virtually unlimited number of possible mixtures to choose from in composing candidate vitrification solutions. It would therefore be of considerable practical utility to have a simple method for predicting the toxicity of a complex mixture of highly concentrated cryoprotectants from first principles.
>
> In the present contribution we show that a simple new compositional variable (qv*) can rationally account, in an organized mammalian tissue, for the toxicity of many complex cryoprotectant mixtures composed to be at total concentrations that are just sufficient to permit vitrification at slow cooling rates at both ambient and elevated pressures. This new compositional variable is proposed to reflect the strength of cryoprotectant hydration within the solution. Based on this interpretation, we were able to predict and successfully test several superior new vitrification solutions with low toxicity for mouse ova, kidney slices, whole rabbit kidneys, and other sensitive systems. These results provide substantial new support for the possibility of developing successful methods for the long-term banking of medically needed tissue and organ replacements.

"Toxicity" in turn is a term that needs visualization if one is to predict it. The visualization settled on by the 21CM researchers depends on the fact that viable mammalian cells contain high concentrations of potassium ions (K^+), versus lower concentrations of sodium ions (Na^+). Though exact proportions vary from experiment to experiment, typically there will be 5 to 7 times as many K^+ ions as Na^+ ions within healthy, untreated cells (controls) [17 p. 30, 19] which can be compared with cells to which cryoprotectants have been applied. In general, the ratio K^+/Na^+ will be lower for these other cases, the difference with controls serving as the indicator of toxicity. One can then search for cryoprotectants that minimize toxicity, that is, yield K^+/Na^+ ratios that are close to the high levels of controls.

Still there is a major problem, as the authors point out. In searching through the many possible CMs for those of low toxicity, testing each mixture individually would be infeasible. Instead the authors propose the quantity qv* which can be calculated for

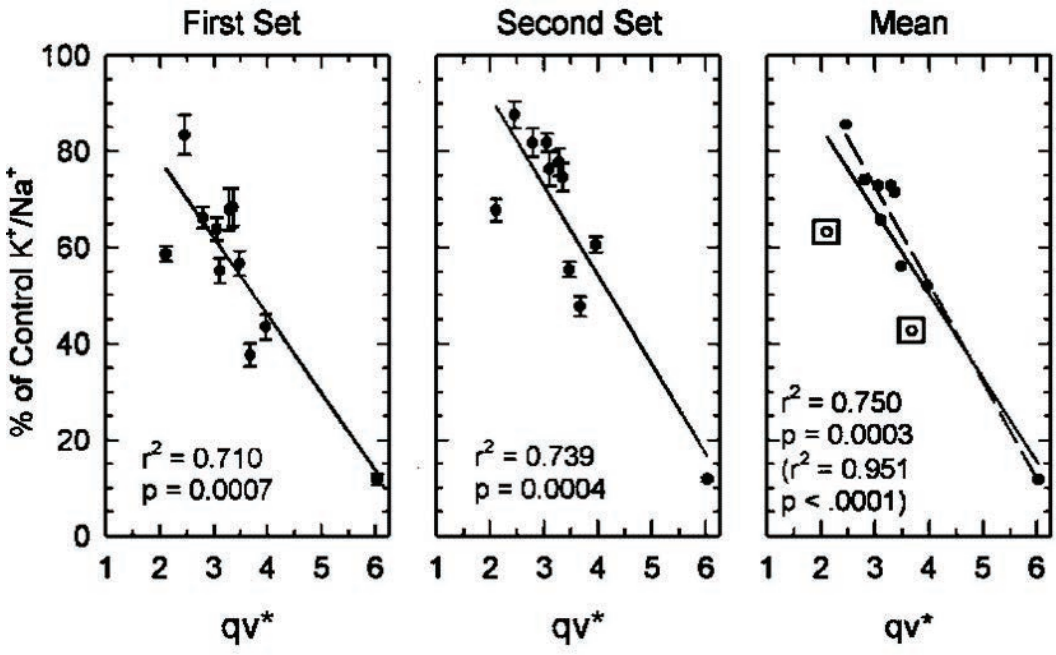

Reproducibility of the correlation between K^+/Na^+ ratio and qv for 11 different vitrification solutions. The left panel recapitulates K^+/Na^+ data collected in Maryland in 1986, while the middle panel shows data collected in California in 2000. The right panel shows the average of these two data sets. The boxed points in the right panel are results that deviated from the regression line by more than 10% in both data sets. Except for these two solutions, the linear regression for the pooled results (dashed line) explains 95% of the variance of the data ($r^2 = 0.95$).*

each mixture and from which the likely toxicity can be estimated. Some results of testing are shown in the chart below, taken from the paper (caption also adapted). Eleven vitrification mixtures or VMs, which are CMs at the minimum concentration needed to vitrify under moderate cooling rates of 10°C/hour, were tested for their toxicity with results plotted against their qv* values. In most cases there was a strong correlation with qv* as shown by the resulting least-squares regression lines: A qv* of 2 or less signified a K^+/Na^+ ratio around 85% or more of that of controls (low toxicity) while a qv* around 6 signified a K^+/Na^+ ratio only 10% that of controls (high toxicity).

Calculating qv* [17,11,10,40]

To calculate qv* we look at a standard volume of a VM, say 1 liter. (The choice of standard volume is for convenience; different choices will give the same results.) This liter of solution will consist of (1) water, (2) "permeating" (or "penetrating") cryoprotectants (PCs) that are needed for vitrification but contribute to toxicity, and (3) other, nonpermeating solutes (NPs) such as carrier solution solutes or ice blockers that are usually present in modest quantity and more or less are inert as far as toxicity is concerned. (Note: by referring to the CM as a "VM" we are assuming it is at the minimum concentration needed to vitrify for a cooling rate of 10°C/hour, as noted above.) qv* is then defined as the ratio M_W/M_{PG} where M_W is the number of moles of water and M_{PG} is the "total polar moles" of the PCs.

A mole of a molecule is the weight in grams equal to its molecular mass. Hydrogen (monatomic, H) has a molecular mass of nearly 1 so 1 mole of hydrogen is about 1 gram. (The number of molecules in a mole is also a fixed quantity, Avogadro's number, about 6×10^{23}.) Water (H_2O) has a molecular mass of 18.015, so this amount in grams makes a mole of water. For a PC we must consider, in addition to how much of the substance we have in moles, the number of "polar groups" it has on its molecule. Polar groups are one of the four chemical groups: amino (NH_2), carbonyl (C=O), hydroxyl (OH), and sulfinyl (S=O) that have been identified as important in the action of PCs. Each PC will have one or more of these polar groups per molecule. So, for example, the cryoprotectant dimethyl sulfoxide (DMSO) has one sulfinyl group per molecule and no other polar groups so its polar groups number is 1. Ethylene glycol (EG) on the other hand has two hydroxyl groups per molecule and no other polar groups so its polar groups number is 2.

We then define the polar moles of a given amount of PC as the number of moles of the PC times its polar groups number. The total polar moles MPG of a VM is the sum of the polar moles of all the PCs in the mixture.

The tabulations below show calculations of qv* for two VMs, VM-1 ("vitrification mixture 1"), and M22. (VM-1 was developed by Dr. Yuri Pichugin and used by the Cryonics Institute. M22 was developed by Drs. Fahy and Wowk at 21CM and is currently used

	A	B	C	D	E	F	G	H	I
1	substance	wt., g	molar mass, g	density, g/ml	vol., ml	moles	polar groups	polar moles	qv*
2									
3	DMSO	300.0	78.13	1.100	272.7	3.840	1	3.840	
4	EG	300.0	62.07	1.113	269.5	4.833	2	9.666	
5	m-RPS-2				20.0				
6	NW				562.2			13.506	
7	W				437.8	24.302			
8	VM-1				1000.0				1.80

Calculation of qv for VM-1. VM-1 consists of an aqueous solution of the two permeating cryoprotectants dimethyl sulfoxide (DMSO) and ethylene glycol (EG) plus the nonpermeating solutes in the carrier solution m-RPS-2. Volumes of these substances that would be present in 1,000 ml of solution are added to obtain the non-water (NW) volume, which is subtracted from 1,000 to obtain the volume of water (W).*

	A	B	C	D	E	F	G	H	I
1	substance	wt., g	molar mass, g	density, g/ml	vol., ml	moles	polar groups	polar moles	qv*
2									
3	DMSO	223.1	78.13	1.100	202.8	2.855	1	2.855	
4	EG	168.4	62.07	1.113	151.3	2.713	2	5.426	
5	F	128.6	45.04	1.133	113.5	2.855	2	5.710	
6	NMF	30.0	59.07	1.011	29.7	0.508	2	1.016	
7	3MP	40.0	106.12	1.114	35.9	0.377	3	1.131	
8	PVP K-12	28.0			28.0				
9	X-1000	10.0			10.0				
10	Z-1000	20.0			20.0				
11	LM5				20.0				
12	NW-LM5	648.1			591.2			16.138	
13	W	411.9		1.000	[411.9]	22.864			
14	M22	1080.0		1.080	1000.0				1.42

Calculation of qv for M22. M22 is an aqueous solution containing the five permeating cryoprotectants dimethyl sulfoxide (DMSO), ethylene glycol (EG), formamide (F), N-methylformamide (NMF), and 3-methoxy-1,2-propanediol (3MP). In addition, it has the three nonpermeating solutes: polyvinyl pyrrolidone (PVP) K-12 and the two ice blockers X-1000, and Z-1000. Besides this are the (nonpermeating) solutes in the carrier solution LM5, here assumed to occupy the same volume as the solutes in the carrier solution m-RPS-2 of VM-1, above. The amount of water (W) is estimated from the known density of M22 rather than by assuming that volumes of all substances add on mixing. Instead, weights of the different components in a liter of solution are subtracted from the weight of the solution, in this case, 1,080 g, to obtain the weight of water (411.9 g) and thus the moles of water (row 13, col. F), which is divided by the total polar moles (row 12, col. H) to obtain qv*.*

by Alcor [7,11,12,13,15,17,18].) Both VMs consist of an aqueous solution containing PCs and one or more NPs.

VM-1, the simpler of the two, has two PCs (DMSO, EG) and one group of NPs, those present in the "carrier solution" m-RPS-2 ("modified renal perfusion solution two"), with the balance water. Volumes in milliliters (ml) of the different components are shown in column E of the table. For the PCs the volume is the weight as shown (col. B) divided by the density (col. D). The carrier solution NPs' combined volume is estimated at about 20 ml/liter [11,10]. The volumes of the PCs and NPs are added to obtain the total non-water volume (NW), which is subtracted from 1 liter = 1,000 ml to obtain the volume of water (W) in the solution, which in turn gives the weight of water in grams (again, 1 ml of water is 1 g). (Note: we assume here that volumes of different substances add on mixing, something that is not strictly true but does appear to hold to a good first approximation in cases like this one [17 p. 29]. Actually, the volumes added together will tend to be slightly less than the sum, due to the way molecules of different liquids "pack" together in solutions. This is taken into account for the more complicated case of M22, below, where it slightly increases the value of qv*.) Knowing the weight and that 1 mole of water weighs 18.015 grams gives amount of water in moles (row 7, column F), and similarly we obtain the amount of the two PCs in moles. Finally, multiplying the moles of a PC by its polar groups number (col. G) gives its polar moles (col. H). The polar moles of the PCs are added (row 6, col. H) and the total divided into the moles of water to obtain the value of qv*, in this case 1.80.

For the case of M22 we have five PCs and multiple NPs. A calculation similar to the above, where it is assumed that volumes of all components add on mixing, is found to give a qv* of 1.34. A more accurate calculation, in which the actual volumetric effects of mixing are taken into account, is shown in the table below; the qv* works out to 1.42. Here we make use of the density of M22, 1.080, so that the amount of water in a liter (1,000 ml; 1,080 g) of solution is the calculations are similar and qv* works out to 1.37, lower than for VM-1, though it appears that both VMs can be considered to have low toxicity in relation to the results shown earlier. For example, rat hippocampal brain slices can be recovered with K/Na ratio's only slightly lower than M22 [10]. More research is needed, both to better assess the severity of damage from existing protocols and to develop better protocols that reduce or eliminate such damage. In any case it appears that the perceptive choice of qv* has accelerated progress in cryoprotectants that might eventually lead to reversible, long-term cryopreservation of tissues and organs, though this goal is still unrealized.

The development of the qv* metric (and the theory of general cryoprotectant toxicity that underpins it) constitutes a major contribution to our understanding of non-specific cryoprotectant toxicity. Further refinements should be possible. For example, as currently calculated, non-penetrating cryoprotectants (like PVP) are considered to be non-toxic, although they do interact with the endothelium and the outer layer of cell membranes. Incorporating the exact hydrogen-bonding strength of the polar groups may further refine qv*. Calculating qv* is only meaningful when cryoprotectant mixtures are chosen with identical vitrification tendencies, but actual toxicity effects can (drastically) change for higher concentrations when a specific threshold is passed to trigger specific toxicity.

Predicting Future Cryonics Case Loads [1,29,30]

Next we consider a very different sort of problem, connected with the fact that cryonics cases are random and unpredictable, so that, for example, relatively long stretches of time can go by between cases, or, on the other hand, there can be several cases over a short time interval. Too many such cases would strain the capabilities of a cryonics organization, so it is desirable to anticipate how often such bunching can be expected.

A main result of some studies by Perry was to develop a formula, based strictly on probability considerations, for the expected waiting time for *n* or more cases to occur over a demand interval of length *t*. Results are shown in the chart below, where curves are plotted for *n* = 2, 3, 4, and 5, assuming a total membership population of 1,000 and an expected average of 7 cases per year, conditions that were fairly approximated at Alcor around 2010-15. (Actually, only the number of cases per year is important, not the member population, but it happens that 7 cases per year is what really occurred for a member population of around 1,000.) We see, for example, that 2 or more cases over a 10-day demand interval *t* could be expected to occur in about a year, whereas (reassuringly) a roughly 30-year wait would be needed for a t of only one day, and 3 or more cases in one day would not be expected for well over 1,000 years. These assumptions cannot be trusted far outside the range of times and other conditions now prevalent and overlook other difficulties, such as pandemics or major accidents which might lead to additional, multiple cases. At any rate, a beginning was made, and some basic results were reassuring.

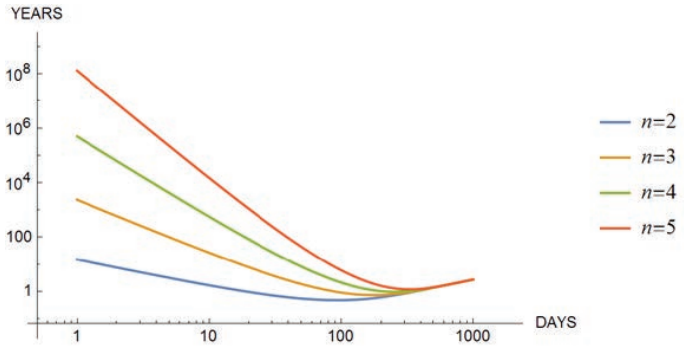

Expected waiting time (years) as a function of demand interval (days) for n *or more cryonics cases, with* n *ranging from 2 through 5, assuming 7 cases per year for member population of 1,000.*

Organizational Stability Issues

One problem not addressed so far, and a very serious one, concerns organizational stability. Some early cryonics organizations were unable to maintain or transfer their patients, which instead were lost (while the organizations themselves subsequently ceased to operate) [28]. Since then much thought and effort has gone into how to prevent the future loss of patients, and particularly, how to ensure that a cryonics organization will be stable and able to maintain its patients indefinitely or until revival is possible. Robert Freitas in a 2010 study addressed this problem for Alcor [35], first noting the financial challenges the organization faced. These included "the current mismatch between cryopreservation funding minimums and actual costs when cryopreservation services are actually rendered, often decades later, combined with Alcor's heavy reliance on bequests and on the continuing generosity of living donors to support core functions."

There had been some "heroic efforts," Freitas noted, to try to address these problems in the past. (Among these was a mathematical study by Perry, noted in passing, estimating how much principal would be needed to fund a cryopreservation indefinitely on interest income, based on different assumptions of rates of return and costs of maintenance [25].) One highly useful tool no one had come up with yet, however, was a "basic top-down econometric model of Alcor's finances." This deficiency Frietas remedies in his study; a footnote explains the field of econometrics and some of its terminology:

> Econometrics is the application of statistical methods to financial or economic data. An econometric model is a set of interlocking mathematical equations that describes the behavior of specific dependent variables when other independent variables fluctuate within their allowed ranges. A "top-down" or "macro" model, as used here, is simplest and can be constructed by starting with top-level aggregated data (e.g., total expenses) and modeling the behavior of that data as a function of several lower-level causative variables (e.g., total number of customers). A "bottom-up" or "micro" model is more complicated (but often more accurate), and may be constructed by starting with low-level unaggregated data (e.g., time series representing actual expenditures on each of 1000 items that the organization must buy in order to provide its services) and then constructing separate sub-models for each of these items, then summing the results of all the sub-models to forecast the top-level aggregated variable(s) of interest. Acquiring the huge amount of data necessary to drive a good bottom-up model can be very expensive and time-consuming.

Due to its complexity, the bottom-up approach that might give the most accurate fit to actual circumstances has been waived in favor of the simpler, top-down approach. Generally, a quantity such as membership, patient total, or expenses will be estimated as a function of time based on linear regression analysis of past data. One example, dealing with Alcor membership, is shown in the figure below; included are projected membership totals up to year 2040 based on various assumptions about growth rates.

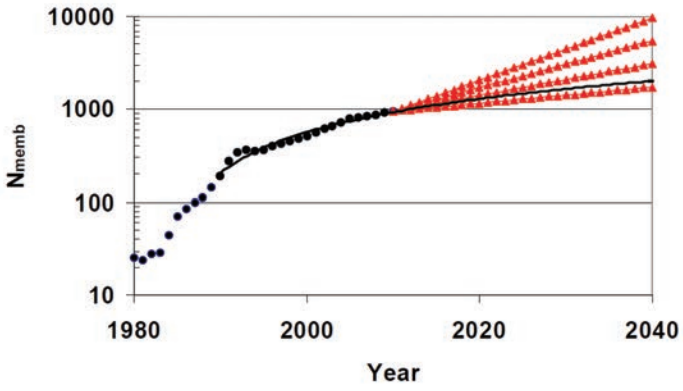

Alcor cryopreservation membership (N_{memb}): actual data, 1980-2009 (black dots); regression formula prediction, 1990-2040 (black curve); N_{memb} prediction for 2010-2040 (red triangles), assuming constant +2%/yr (bottom), +4%/yr, +6%/yr, or +8%/yr (top) growth rates.

In a follow-up article in *Cryonics* Freitas summarizes what has been accomplished in his study [36]:

> The analysis starts by creating a model of Alcor's expenses using historical data from 1990-2008. Statistical correlation is employed to predict the expense data using three independent variables: number of members, number of cryopatients, and number of cryopreservations per year. Using various assumed growth rate scenarios for these three independent variables, Alcor expenses can be projected forward 30 years into the future. The analysis continues with the creation of a similar model of Alcor's revenues based on historical data from 1990-2008. Statistical correlation is again employed to predict the revenue data using sub-models for each of Alcor's five principal consolidated revenue sources: (1) dues, (2) standby fees, (3) proceeds from cryopreservations, (4) Patient Care Trust (PCT) earnings, and (5) grants, donations and bequests. Each revenue stream can be predicted using the same three independent variables as before. This allows Alcor's revenues – and, after subtracting predicted expenses, any budget shortfalls or surpluses – to be projected forward 30 years into the future.

Serious concerns are raised about Alcor's then-current fiscal policies. Freitas then offers some recommendations for possible funding options, involving dues increases, cost-of-living adjustments, or donations that would "produce long-term budgetary stability at Alcor." He feels that "members should

always have available to them at least one viable option that includes the permanent grandfathering of their account." Some detailed recommendations for further studies are also presented.

Freitas's work served as an important reference in later discussions and policy decisions on funding issues at Alcor – a lengthy story which can only be hinted at here [2,3].

Cryonics and Machine Learning

Up to now we have considered several applications of basically straightforward numerical mathematics to problems in cryonics. Here instead we will delve into machine learning, a branch of artificial intelligence which considers how a computer algorithm might be modified to improve its performance on tasks of the type that are considered to require intelligence if performed by humans. An example would be a game-playing program where we are interested in the process of generating successively better versions of the program so that the machine "learns" to play better. The learning process could involve human guidance ("supervised learning") or be entirely automated ("unsupervised learning").

Attempts to apply machine learning to cryonics problems include an early study by Perry in which EM images of rat brain cortical tissue, initially represented as 2-D numerical arrays, were collapsed into 1-dimensional "signatures" using Fourier analysis. A primitive machine learning technique was used to find a vector whose inner product with the signature would approximate the amount of ischemic time of the corresponding brain tissue. Results were informally published in 2011 [27]. The intent of the work was to make some headway toward understanding what changes occur in brain tissue as a result of postmortem ischemia. In this way light might be shed on the problem of whether and how cryonics patients might be revived at a future date. Only limited success was achieved, precluding such an ambitious goal, but the new approach opened possibilities for future work.

Recently a second study has been completed [9], that uses a far more sophisticated machine learning technique implemented by Michael Maire to distinguish tissues with different ischemic exposures. The study is important in part for its application of what can be considered state-of-the art deep learning to a problem of cryonics. It is also important for its detailed, systematic study of changes that occur in brain tissue as a function of postmortem ischemic exposure, something that is especially of interest in cryonics but not much pursued outside the field. Perhaps, then, we are closer to the goal of using computational methods to show how important problems connected with cryonics revival could be solved, though much remains to be done before any substantial success can be claimed in achieving such a goal.

For the machine learning application, a deep convolutional neural network is used to addresses a classification task involving electron microscopy (EM) images of rat cortical brain slices (tissue samples). The tissue samples are in two main groups, one exposed to warm ischemia (rat normothermic body temperature, 37°C, similar to human body temperature) the other to cold ischemia (0°C). For the warm ischemia the exposure times range from 0 (control) up to 81 hours. For the cold ischemia the (overlapping, but generally much longer) times are from 0 to 6 months (4,464 hours). (Note that in view of the Q-10 rule we expect a much slower rate of tissue degradation per unit of time at the lower temperature.) A program known as the "computational model" is to be given data derived from an EM image of a tissue sample and return information pertinent to the ischemic history of that sample, that is to say, what temperature the sample was exposed to and for how long.

It turns out that to develop or "train" a computational model that simply outputs this information, and produces reliable answers, would be a very ambitious undertaking at this point. (Indeed, it would require considerable, specialized expertise for a human to carry out such a task, where possible at all. Some of the cases of different ischemic times were hard to distinguish; see below.) So instead a simpler task was chosen for which, initially, the sample images were sorted into batches, with each batch consisting of images of samples which were all alike in times and temperatures. So, one batch might represent controls with no ischemic time at all, another with 1 week (168 hours) of cold exposure (0°C), and so on. Altogether there were n batches of images. (In the results as shown below, $n = 19$, with about 14,000 images per batch. Each image, consisting of a "patch" cropped and downsampled from a larger image, was a 128 x 128 array of pixels showing brain structure, with pixel width about 13 nanometers. Thickness of brain tissue slices as imaged was about 70 nanometers.) To test a given computational model an image is chosen from batch m, say, together with a pair of batch numbers (p, q), in which one of the p, q, is m, the "right answer" and the other is the wrong answer. The computational model is then tasked with making the right choice of batch number among the two possibilities. The performance of the model can be gauged by testing it on a great many examples of images covering all the different combinations of p and q and recording the percentage of correct answers it gives for each combination.

The problem then becomes one of arriving at a computational model that produces correct answers as often as possible. This is the "training" problem. For this we start with a "baseline model" that is totally ignorant of how to answer and gives answers at random, thus achieving 50% correctness (chance only) for each case. During the training phase, when the model gives an incorrect answer a correction is applied to the model to try to improve its performance, and the model in the end should perform well on the training examples. It is then tested on similar examples it was not trained on, images it has not seen before, for each combination of p, q, to see if it will perform well, or at least better than chance, on these cases too. Success

means it has acquired "understanding" of the problem at hand and can impart useful insight.

Additional details of the process are here excerpted; more will be found in the original paper:

> Deep convolutional neural networks have emerged as the standard method for image classification. They are composed of many stacked layers, each consisting of a filter set followed by a nonlinearity and acting on the output of the previous layer. Given an image as input, filters in the first layer typically extract basic structures such as edges, whereas those in deeper layers build more abstract features. Trained on a large dataset of natural images, their learned representations are transferable to a variety of image interpretation tasks. An existing network consisting of 16 convolutional layers trained on natural images was used as a module for generating a descriptor for any generic input image. This network outputs a descriptor in the form of a 4096-dimensional vector produced by its deepest layer.
>
> Within this computational framework, we explored the question of the discriminability of electron microscopy images of samples subjected to different ischemic time and temperature conditions. … Tested on an equal number of images from each of … two conditions [that is, (*p*,*q*) where first *p* then *q* is the right answer], the baseline model (random chance) would have an accuracy of 50%. The extent to which a trained model is capable of outperforming chance is indicative of the visual distinctiveness of neural tissue subject to the two different experimental conditions. All possible pairs of ischemic conditions were considered, and a discriminative model was trained for each pair.

Principal results are shown in the chart below, for nineteen batches of images covering both the cold and warm ischemia cases as noted above, each batch label beginning either "c" ("cold") or "w" ("warm") followed by number of hours. We start with the outlier case of cold ischemia (c4644, batch 1), count downward through the time values to c0 (batch 9), then upward through the warm time values to w81 (batch 19). For each (*p*,*q*) pair of possible batches we show the color-coded percentage of correct answers of the trained computational model when first *p* then *q* is the right answer. Thus the overall result for (*p*,*q*) will equal that for (*q*,*p*); in the chart the results are shown only for the values above chance with neutral fill-in (50%, chance level) otherwise, and in the trivial case of *p*=*q*.

Among the interesting details is that, as expected, ischemic times that are close together with similar temperatures are generally not as strongly distinguished as cases with greater disparity, some nearby cases approaching or (in the chosen representation) equaling chance only. Another interesting property is that most of

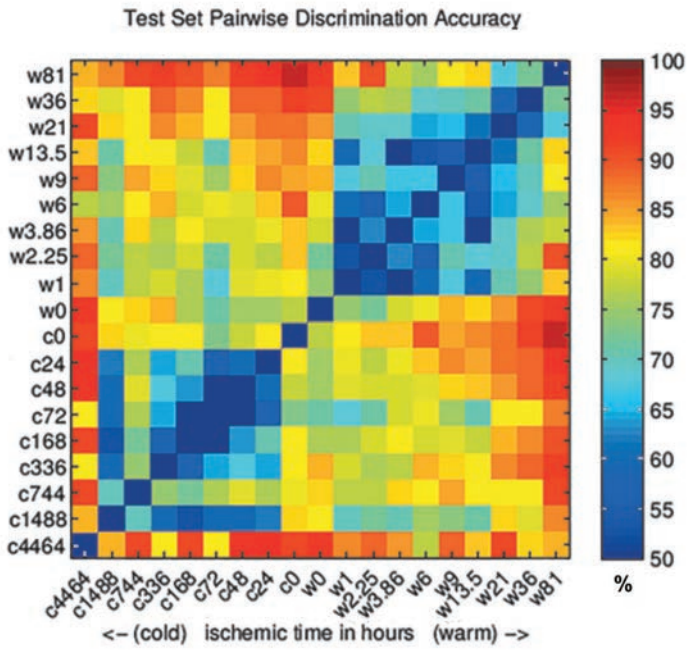

the cold ischemia cases are rather strongly distinguished from the warm cases (upper-left and lower-right quadrants), regardless of exposure times, and better distinguished than cases within each group are from each other (lower-left and upper-right quadrants). It appears that rather different types of distinguishable damage occur at the two different temperatures. A third observation is that controls tend to be especially well-distinguished from all non-controls, suggesting that some important damage occurs early on; for example, in the case of warm ischemia, with E-HIT of 1 in the first hour. The ability of the algorithm to successfully distinguish between the very early stages of ischemia and the late (necrotic) stages of ischemia appears to support the practice of cryonics, which takes advantage of the fact that there is an intermediate "island of ultrastructural stability" (lasting for hours at normal body temperature) which permits preservation of identity-critical information. Overall, then, the work may be an opening wedge that could lead to greater application of machine learning and/or more general artificial intelligence in cryonics problems, with exciting conclusions.

Closing Remarks

In cryonics we are concerned with whether our procedures will be good enough for the hoped-for revival of our patients someday. We must do the best we can without knowing the answer to this all-important question. The driving interest we have in ensuring our procedures are "good enough" without knowing at what point we may have succeeded impels us to examine and analyze our procedures closely, and it might be argued that this is mainly where mathematics comes into play.

We have seen how mathematics was used decades ago to help understand and address problems of cryoprotective perfusion

and heat flow. Some years later the problem of ischemic injury prior to cryogenic cooling came under scrutiny, with a metric, the E-HIT, proposed to estimate the amount of injury resulting from different possible cryopreservation protocols. Meanwhile some studies were done that modelled cryoprotectant toxicity to arrive at better vitrification solutions for cryogenic cooling. A study was done to estimate the likelihood of excessive cryonics caseloads based on an assumption that cases occur at random. Alcor finances came under scrutiny in another study, with long-term projections of expenses and revenues under different assumptions of funding requirements and likely caseloads. Finally, there is the recent study in which an attempt is made to apply machine learning to gain a computerized understanding of the progress of ischemic damage before sufficient cooling can take place.

Overall, there is much that remains to be done. There is also much that *has* been done of a more routine nature, connected with the automation of perfusion and cooldown procedures, for example, that was not reported here (notwithstanding the considerable and diligent efforts that were involved). Many of the mathematical studies reported here do not appear to have had any profound effect on policies and protocols to date. Some notable exceptions were the studies on cryoprotectants and the studies about the long-term effects of "grandfathering." This situation will hopefully improve in the future when we may expect that deeper mathematical studies can be conducted, again possibly involving artificial intelligence. Meanwhile we can look for relatively simple ways to improve our protocols; further development of E-HIT metric and combining them with CT scan results might be one good place to start. ■

The authors thank Greg Fahy and Brian Wowk for assistance, particularly with the section on cryoprotectants Hugh Hixon also provided some helpful comments on current cryonics procedures regarding early-stage cooling.

Image Credits

Images relating to heat flow and perfusion modeling were taken from the references cited and sources should be clear. Plots of "% of control K^+/Na^+ vs. qv*" were based on graphics from [17]. The graphic showing Alcor membership and projected membership over time was adapted with permission from the reference cited ([35]); "Test Set Pairwise Discrimination Accuracy" was based on graphics from [9]; artwork by RMP unless otherwise noted.

Endnotes

1. Aaron Drake, "Is there a Need for More than One Standby Team?" *Cryonics* 39(2) (Mar.-Apr. 2018) 10-11.
2. Alcor Management and Board, "Cryopreservation Funding and Inflation: The Need for Action" (30 Sep. 2011), https://alcor.org/Library/html/CryopreservationFundingAndInflation.html, accessed 22 May 2020.
3. (Alcor financial questions) https://www.alcor.org/FAQs/faq05.html, accessed 22 May 2020.
4. Art Quaife, "Mathematical Models of Perfusion Processes," *Manrise Technical Review* 2(2-3) (Mar.-Jun. 1972), 28-75.
5. "New Perfusion Studies by Quaife," *The Outlook* 3(4) (Apr. 1972) 1.
6. Art Quaife, "Heat Flow in the Cryonic Suspension of Humans: A Survey of the General Theory," *Cryonics* 6(9) (Sep. 1985) 9-30.
7. Aschwin de Wolf, "Vitrification Agents in Cryonics: M22" (8 Jul. 2008), https://www.biostasis.com/vitrification-agents-in-cryonics-m22/, accessed 6 May 2020.
8. Aschwin de Wolf, "Critical Cooling Rate to Prevent Ischemic Brain Injury," (10 Jul. 2008), https://www.biostasis.com/critical-cooling-rate-to-prevent-ischemic-brain-injury/, accessed 23 Apr. 2020.
9. Aschwin de Wolf, Chana Phaedra, R. Michael Perry, and Michael Maire, "Ultrastructural Characterization of Prolonged Normothermic and Cold Cerebral Ischemia in the Adult Rat," *Rejuvenation*, Published Online:12 Mar 2020, https://doi.org/10.1089/rej.2019.2225, accessed 20 Apr. 2020 (to appear in print about Jun. 2020).
10. Aschwin de Wolf, "Vitrification Agents in Cryonics: VM-1" (31 Mar. 2008), https://www.biostasis.com/vitrification-agents-in-cryonics-vm-1/, accessed 26 May 2020.
11. Ben Best, "Viability, Cryoprotectant Toxicity and Chilling Injury in Cryonics," https://www.benbest.com/cryonics/viable.html#toxicity, accessed 25 Apr. 2020.
12. Ben Best, "A History of Cryonics," https://www.benbest.com/cryonics/history.html, accessed 6 May 2020.
13. Brian Wowk, "How Cryoprotectants Work," *Cryonics* 28(3) (3 Q 2007) 3-7, https://alcor.org/Library/pdfs/How-Cryoprotectants-Work.pdf, accessed 5 May 2020.
14. Brian Wowk, "Thermodynamic Aspects of Vitrification," Cryobiology 60(2010) 11-22.
15. "CI-VM-1 Cryoprotectant and CI-Carrier Solution Used for Vitrification," http://www.cryonics.org/research/CI-VM-1.html, accessed 5 May 2020 via Wayback.

16. (general article on cryonics). https://en.wikipedia.org/wiki/Cryonics, accessed 19 Apr. 2020.

17. Gregory M. Fahy, Brian Wowk, Jun Wu, and Sharon Paynter, "Improved Vitrification Solutions Based on the Predictability of Vitrification Solution Toxicity," *Cryobiology* 48 (2004) 22–35, http://www.21cm.com/pdfs/improved_vitrification.pdf, accessed 25 Apr. 2020.

18. "Advantageous Carrier Solution for Vitrifiable Concentrations of Cryoprotectants, and Compatible Cryoprotectant Mixtures," 21st Century Medicine, Inc.. 27 Jul., 2001, https://patents.justia.com/patent/6869757, accessed 5 May 2020.

19. Gregory M. Fahy, private communications 27 Apr.-21 May 2020 (emails).

20. Steven B. Harris, Michael G. Darwin, Sandra R. Russell, Joan M. O'Farrell, Mike Fletcher, Brian Wowk, "Rapid (0.5°C/min) Minimally Invasive Induction of Hypothermia Using Cold Perfluorochemical Lung Lavage in Dogs," *Resuscitation* 50 (2001) 189-204.

21. Hugh Hixon, "How Cold Is Cold Enough?," orig. "The Question Column," *Cryonics* 6(1) (Jan. 1985) 19-25, https://alcor.org/Library/html/HowColdIsColdEnough.html#:~:text=Below%20%2D100%C2%B0C%2C%20the,move%20and%20undergo%20chemical%20change., accessed 20 Jun. 2020.

22. Michael Darwin, "A History of DMSO and Glycerol in Cryonics," *Cryonics* 28(3) (3Q 2007) 8-11.

23. https://www.alcor.org/blog/official-alcor-statement-concerning-marvin-minsky/, accessed 19 Apr. 2020.

24. R. Michael Perry, "Mathematical Analysis of Recirculating Perfusion Systems, with Application to Cryonic Suspension," *Cryonics* 9(10) (Oct. 1988), 24-38.

25. R. Michael Perry, "A Mathematical Model of Long-Term Storage Costs," Cryonics 11(8) (Aug. 1990), 37-38.

26. R. Michael Perry, "Toward a Measure of Ischemic Exposure," *Cryonics* 17(2) (2Q 1996), 21.

27. R. Michael Perry, "Algorithmic Estimation of Cortical Autolysis," https://alcor.org/Library/pdfs/Algorithmic_Estimation_of_Cortical_Autolysis.pdf) (2011), accessed 19 Apr. 2020.

28. R. Michael Perry, "Suspension Failures: Lessons from the Early Years," https://alcor.org/Library/html/suspensionfailures.html, accessed 22 May 2020.

29. R. Michael Perry,"Estimating and Forecasting Alcor Resource Requirements: Are Cases Random?" *Cryonics* 36(12) (Dec. 2015) 22-24.

30. R. Michael Perry, "Expected Wait Times for Multiple Cryonics Cases in a Specified Time Interval" (7 Mar. 2016), https://mega.nz/file/MBtmQCDR#vZ9C4ReUJwpnMOFJt1x7II0Q52edIL_QEy5wEZNHLMA, accessed 6 May 2020.

31. R. Michael Perry, "Cryonics Newsletters: Some Historical Highlights," *Cryonics* 39(1) (Jan.-Feb. 2018) 30-39.

32. Mike Darwin, "Longtime Alcor Member Enters Biostasis," *Cryonics* 9(5) (Jun. 1988) 2-13, https://alcor.org/Library/html/casereport8806.html, accessed 19 Apr. 2020.

33. https://alcor.org/Library/pdfs/casereportA1108RobertBinkowski.pdf, accessed 7 Apr. 2020.

34. Robert C. W. Ettinger, *The Prospect of Immortality*, Doubleday, 1964, 106-11.

35. Robert A. Freitas Jr., "Scenario Analysis using a Simple Econometric Model of Alcor Finances," https://www.alcor.org/Library/pdfs/EconometricModelOfAlcorFinances.pdf, summary at https://www.semanticscholar.org/paper/Scenario-Analysis-using-a-Simple-Econometric-Model-Freitas/f2fb11480df1c7adccad238d787bfd92bf0bf234, both accessed 10 May 2020.

36. Robert A. Freitas Jr., "Long-Term Financial Stability in Cryonics," *Cryonics* 31(3) (3Q 2010) 4, https://alcor.org/Library/html/FinancialStability.html, accessed 24 May 2020.

37. https://www.reddit.com/r/Futurology/comments/4dj8v9/ray_kurzweil_remembers_marvin_minsky_confirms/, accessed 27 Mar. 2020.

38. Kazunori Tahara, Takashi Murakami, Jun Fujishiro, Masafumi Takahashi, Seiichiro Inoue, Kohei Hashizume, Kenjiro Matsuno, and Eiji Kobayashi, "Regeneration of the Rat Neonatal Intestine in Transplantation," Ann Surg. 2005 Jul; 242(1): 124–132. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1357713/, Accessed 25 Apr. 2020.

39. Steven B. Harris, "Initial Cooling in Cryonics," unpublished manuscript (22 Feb. 2003).

40. (substances). 3-methoxy-1,2-propanediol (3MP): https://www.sigmaaldrich.com/catalog/product/aldrich/260401?lang=en&region=US; dimethyl sulfoxide (DMSO): https://en.wikipedia.org/wiki/Dimethyl_sulfoxide; ethylene glycol (EG): https://en.wikipedia.org/wiki/Ethylene_glycol; formamide (F): https://en.wikipedia.org/wiki/Formamide; N-methylformamide (NMF): https://en.wikipedia.org/wiki/*N*-Methylformamide; polyvinyl pyrrolidone (PVP): https://www.ncbi.nlm.nih.gov/pubmed/19839742. All accessed 3 May 2020.