

1st Danish – Norwegian Study Group with Industry (ESGI 137)

Ålesund, June 11-15, 2018

Poul Hjorth, Dietmar Hömberg, and
Siebe van Albada (eds.)



© NTNU

NO-7491 TRONDHEIM, NORWAY

FIRST EDITION: NOVEMBER 2018



Contents

Preface	5
1 Induction tempering	7
1.1 Background and problem statement	7
1.1.1 Problem statement	9
1.1.2 Additional prerequisites	9
1.1.3 Tasks to be done	10
1.2 Discussion of tempering parameters	11
1.3 Optimization problem	12
1.3.1 Optimal control problem	12
1.3.2 Approaches for the control $u(t)$	13
1.4 Numerical Solution of the Heat Equation	14
1.4.1 Spatial discretisation	14
1.4.2 Time discretization	15
1.5 Results	15
1.5.1 Pulsed Heating	15
1.5.2 Exponential Heating	17
1.5.3 Linear Splines	17
2 Total Hip Arthroplasty	21
2.1 Problem definition	21
2.2 X-ray scaling using a reference sphere	26
2.2.1 Automatic locating and sizing of the sphere	26

2.2.2	Scale error caused by sphere location	26
2.2.3	Estimation of the elevation above the X-ray plate	29
2.3	Locating and measuring femurs	31
2.3.1	Automatic location and orientation	31
2.3.2	Estimation of channel widths	31
2.3.3	Automatic location of minor trochanters	33
2.4	Locating and measuring the femoral heads	35
2.5	Identifying pelvic symmetry	37
2.6	Estimating 3D shape from a 2D X-ray	38
2.6.1	3D to 2D projection	41
2.6.2	2D to 3D reconstruction	43
2.6.3	2D/3D hybrid templating	45
2.7	Conclusions and Future Work	45
3	Predicting salmon weight	50
3.1	Introduction	50
3.2	Deterministic analysis	51
3.3	Statistical analysis	52
3.3.1	Preliminary analysis	52
3.3.2	Sampling consequences	52
3.3.3	The effect of having two populations	53
3.3.4	Statistical model	54
3.3.5	Simulation study	55
3.4	Summary	56
3.5	Appendix	58
4	Electrode paste softening and flow	63
4.1	The Challenge	63
4.2	Previous Work	66
4.3	Early Developments	67
4.3.1	Heat Transport	67
4.3.2	Lubrication Flow	68
4.4	A Fully Coupled Model	69
4.4.1	Conceptual Model	69
4.4.2	Equations	70
4.4.3	Boundary conditions	71
4.4.4	Non-dimensionalisation	71
4.4.5	Three regions	71
	List of Participants	75



Preface

European Study Group with Industry (ESGI) is Europe's leading workshop bringing together mathematicians and industrial companies. These week-long workshops have been held annually since 1968 and were previously known as the Oxford Study Groups with Industry. These events are now spreading across Europe (indeed to other continents) where they attract leading mathematicians to work on industrial problems.

The first ESGI in Norway was organized by mathematicians from the Norwegian University of Science and Technology supported by colleagues from the Technical University of Denmark in Lyngby. It took place at the NTNU campus Ålesund, June 11-15, 2018.

More than 30 participants from six European countries, Saudi Arabia, and New Zealand worked on 4 challenging problems.

The first one was posed by the company EFD Induction, a Norwegian producer of induction heating equipment. They asked for an optimal strategy for the induction tempering process such that an acceptable hardness gradient in the hardened steel parts is accomplished.

Problem 2 came from Stavanger University Hospital. They asked us for an analysis of options to improve safety for patients with total hip arthroplasty by improving preoperative digital templating strategies.

The third problem was suggested by SalMar, a Norwegian producer of farmed salmon. Here, the goal was to predict the average weight of the fish during production in sea phase, the slaughter weight and weight distribution analysing also the measurement uncertainties.

Last but not least, Elkem, a supplier of silicon-based advanced materials asked us to develop a dynamic mechanical model of the different parts of an electrode during its movement in a smelting furnace.

The problems were taken up and discussed by an enthusiastic group of participants. They came up with interesting results presented to the problem posers on Friday morning in the usual study group style.

The present report summarizes the results of this study group. We envisage that this will pave the way to further collaborations between academia and industry. All in all the event was so well-received by all participants that we hope to be able to organise another study group in Norway in two years.

We would like to express our gratitude to our sponsors, NTNU and the COST Action TD1409 Mathematics for industry network for their financial support. Our thanks go to Anniken Th. Karlsen, head of Department of ICT and Natural Sciences at NTNU Ålesund for her encouragement and support of the event. Finally, we would like to thank Mark McGuinness for the beautiful pictures taken during the study group.

Trondheim, 21 November 2018

Poul Hjorth, Dietmar Hömberg, and Siebe van Albada



Problem 1. Induction tempering

PROBLEM PRESENTED BY: D. IVANOV, EFD INDUCTION

STUDY GROUP: H. HANCHE-OLSEN, T. PETZOLD, I. ROPER, T. ROY, C. VEJE

1.1 Background and problem statement

The problem posed by the company EFD Induction deals with the process of induction tempering.

Tempering is part of the heat treatment for workpieces made of steel. During the hardening of the steel, the phase martensite has formed, which has high hardness but is very brittle. In order to improve the properties of the workpiece, tempering is performed, where the part is heated up again to a defined temperature. During this additional heating, some of the internal stresses in the martensitic phase are released. Regarding the microstructure of the steel, carbides are formed during the heating that lead to a removal of the excess hardness and an overall improvement of the toughness of the steel.

Therefore, tempering results in an increase of ductility at the cost of a reduction of hardness. Based on the desired properties of the final product, the tempering process needs to be controlled to achieve the required properties. The parameters to be controlled are the temperature and the heating time.

One way of tempering is to use a furnace, where the parts are heated at constant temperature for a specified time duration. The usual heating times are in the range of one hour. On the other hand, the tempering can be done using inductive heating. The advantage of induction tempering is that the heat is generated

directly in the part. As a consequence, the heating times are much shorter, usually in the range of 10 to 30 seconds, see Fig. 1.1, which makes the process very energy efficient.




Figure 1.1: Comparison of heating time and temperature for tempering in a furnace and induction tempering.

Due to the skin effect, the heating of the part using induction is not homogeneous. More heat is generated close to the surface of the part, which results in higher temperatures at the surface. Therefore, the hardness reduction during the tempering is not as homogeneous as for tempering in a furnace. The hardness of a part after the tempering is exemplary shown in Fig. 1.2.




Figure 1.2: Comparison of the hardness after the tempering.

1.1.1 Problem statement

The aim of the study group was to investigate the effect of the hardness reduction during the process of induction tempering. In particular, the aim was to find heating strategies to reduce the inhomogeneity of the hardness gradient.

During the tempering, the power of the inductive heater can be controlled. Instead of providing a constant power over a certain time, the power level and the heating time can be adjusted to control the temperature profile and consequently the hardening reduction in the workpiece. One way of adjusting the power input can be a pulsed heating. But also different strategies with varying input power over time are possible.

The problem results in an optimization problem to find an optimal heating strategy to achieve a desired tempering effect. In particular, the following requirements were set by EFD:

- Heating/tempering should be achieved in *shortest time*.
- The properties of the final product in terms of a desired hardness/hardness reduction are predefined.
- The gradient of the hardness reduction needs to be minimized.
- Usually, an acceptable hardness gradient is defined by the customer, there is no need to do better than desired.

1.1.2 Additional prerequisites

Regarding the mathematical model, inductive heating can be described by a set of partial differential equations. The electro-magnetical effects are modelled by Maxwell's equations. The temperature distribution can be obtained from the heat equation.

The tempering effect can be measured in terms of the hardness reduction, which in turn is related to a phenomenological parameter P that depends on temperature (influence of the heating temperature) and time (influence of heating duration). One example for P is given in terms of the *Hollomon-Jaffe* parameter ([HJ45](#)), which in the isothermal case is given by

$$P = T(C + \log t). \quad (1.1)$$

With respect to the geometry, an infinitely long cylindrical workpiece is considered in this work. Therefore, one can make use of the rotational symmetry, see Fig. 1.3.

With these assumptions, it is possible to derive an analytical solution for Maxwell's equation. For the induced current density there holds

$$J(r) = -\frac{\text{ber}'\left(\frac{\sqrt{2}r}{\delta}\right) + j\text{bei}'\left(\frac{\sqrt{2}r}{\delta}\right)}{\text{ber}\left(\frac{\sqrt{2}R}{\delta}\right) + j\text{bei}\left(\frac{\sqrt{2}R}{\delta}\right)} \frac{\sqrt{2}H_s}{\delta}, \quad (1.2)$$




Figure 1.3: Workpiece geometry

where $\delta = \sqrt{\frac{2}{\sigma\omega\mu}}$ denotes the material dependent penetration depth of the induced current with μ the permeability, σ the electrical conductivity and $\omega = 2\pi f$ the radial frequency; H_s denotes the surface magnetic field intensity and $\text{ber}(\cdot)$ and $\text{bei}(\cdot)$ denote certain Bessel functions, respectively Kelvin functions.

The heat equation in the rotational symmetric setting is written as

$$c\rho \frac{\partial T}{\partial t} = \frac{\kappa}{r} \frac{\partial}{\partial r} \left(r \frac{\partial T}{\partial r} \right) + u(t)P_v(r) \quad \text{in } [0, R], \quad (1.3)$$

where c denotes heat capacity, ρ the density, κ the heat conductivity and $P_v(r) = \frac{1}{\sigma} J(r)^2$ the heat power. As boundary condition, there holds a zero flux condition at $r = 0$ and at the outer boundary $r = R$ we impose a cooling law

$$-\kappa \frac{\partial T}{\partial r} \Big|_{r=R} = \alpha(T - T_\infty) \quad (1.4)$$

with α the heat transfer coefficient and T_∞ the ambient temperature. Furthermore, a suitable initial condition for the temperature is defined: $T(r, 0) = T_0$.

The function $u(t) : [0, t^*] \rightarrow U$ with U a suitable space, serves as a control. It could represent a normalized current intensity or the relative power of the inverter.

1.1.3 Tasks to be done

- In order to solve an optimization problem, it is necessary to solve the state equation, i.e. to determine the temperature for any given control $u(t)$. It has to be investigated, if it is possible to obtain an analytical solution for any given control $u(t)$. Otherwise, the temperature has to be determined by numerical methods, e.g. finite differences or finite element methods.

- The parameter P indicating the tempering effect, equation (1.1), is related to isothermal situations. In order to apply it to induction tempering, a suitable generalization of P for non-isothermal situations needs to be found.
- A constrained optimization problem needs to be defined, consisting of a cost functional and suitable constraints on the control $u(t)$ and the state variables such as temperature and the parameter P .
- In order to obtain the heating strategy in terms of the control $u(t)$, the optimization problem has to be solved using suitable algorithms. Different approaches for the representation of the control $u(t)$ need to be discussed.

1.2 Discussion of tempering parameters

In order to quantify the tempering effect, a parameter P is introduced to phenomenologically describe the effect of tempering. This parameter is then related to a hardness reduction.

One possibility to define this parameter P is the Hollomon-Jaffe parameter (HJ45), see e.g. equation (1.1). It relates the effect of heating temperature and heating time to the hardness reduction. From equation (1.1) one can deduce, that the temperature is very important. Due to the logarithmic dependence on time, an increase of the heating time has a smaller effect than an increase in temperature.

In literature, a variety of parameters exist. During the workshop, the *Dorn-parameter*, (RGT87), has also been discussed in detail. It is given by

$$P = \exp\left(-\frac{Q}{RT}\right)t, \quad (1.5)$$

where Q denotes an activation energy and R the gas constant. While during tempering in a furnace, the temperature is constant over a long period of time, in induction tempering, one has to deal with a non-isothermal situation. Therefore, a suitable representation of tempering parameters for the non-isothermal case is necessary. This is for example considered in (RFG+10).

The authors describe a general framework for the extension of isothermal time-temperature parameters to non-isothermal conditions. They conclude that the Hollomon-Jaffe parameter is not suitable to give a non-isothermal representation. They claim that a generalized Dorn parameter in the form

$$P(t, T) = \int_0^{t^*} \exp\left(-\frac{Q}{RT}\right) dt \quad (1.6)$$

is “regarded to be the sole non-isothermal complex parameter to which a rigorous, physically well-founded interpretation is attributed”, (RFG+10). Therefore, in this work, the generalized Dorn parameter will be considered.

1.3 Optimization problem

In this section, we will formulate the optimization problem for the induction tempering problem. It is the aim to compute a heating strategy, i.e. a realization of the control $u(t)$, to achieve the desired tempering in shortest time. Furthermore, the gradients of the hardness reduction in the surface region of the workpiece need to be minimized. Usually, an acceptable hardness gradient is defined by customers, which will be incorporated in the optimization problem.

1.3.1 Optimal control problem

As mentioned before, it is the aim to find the time minimal heating strategy to the induction tempering problem in terms of the control $u : [0, t^*] \rightarrow U$, that maps from the time interval $[0, t^*]$ with variable end time t^* onto some space U . The space U could be the unit interval $[0, 1]$ in case we consider a normalized current intensity. We consider the following cost functional

$$\min_{u: [0, t^*] \rightarrow U} t^*,$$

where the control has to satisfy e.g. the following control constraints

$$0 \leq u(t) \leq 1 \quad \text{for all times } t \in [0, t^*].$$

The temperature $T(r, t)$ of the cylindrical workpiece and the tempering parameter $P(r, t)$ define the state of the system. These quantities have to satisfy the following state equations

$$\begin{aligned} c\rho \frac{\partial T}{\partial t} &= \frac{\kappa}{r} \frac{\partial}{\partial r} \left(r \frac{\partial T}{\partial r} \right) + u(t)P_v(r), & T(0) &= T_0 \\ \frac{\partial P}{\partial t} &= f(t, T), & P(0) &= 0 \end{aligned}$$

together with the boundary condition (1.4). From the problem description, we infer the following state constraints on the parameter P :

$$P(t^*, R) = P_{\text{desired}} \tag{1.7}$$

$$|P(t^*, R) - P(t^*, R_{HD})| \leq \varepsilon. \tag{1.8}$$

Here, R denotes the outer radius of the workpiece, P_{desired} is a given desired tempering effect (at the surface), R_{HD} denotes the radius related to the hardening depth and ε is a given tolerance for the gradient of P . One could also think of additional state constraints with respect to the temperature, e.g. a maximum temperature in order to avoid overheating or a restriction to temperature gradients in order to prevent additional thermal strains in the workpiece.

In this report, only the constraints with respect to P are considered. In the optimization, the inequality constraint $g(P) := |P(t^*, R) - P(t^*, R_{HD})| - \varepsilon \leq 0$ will be considered using a penalty function. A possible choice is the quadratic penalty function

$$s(g(P)) = \beta \max\{0, g(P)\}^2$$

with a penalty parameter β . In case the constraint is violated, i.e. $g(P) > 0$, then a positive weight is added to the cost.

Since we are interested in minimizing the heating time, a condition to define the end time, i.e. a stopping condition, must be defined. In this work, we stop the heating as soon as a desired tempering effect at the surface is achieved, which is described in terms of the parameter P . Therefore, the constraint $P(t^*, R) = P_{\text{desired}}$ is used as stopping criterion.

In order to summarize, the optimization problem is the following:

$$\min_{u: [0, t^*] \rightarrow U} t^* + \beta \max\{0, |P(t^*, R) - P(t^*, R_{HD})| - \varepsilon\}^2 \quad (1.9)$$

such that

$$c\rho \frac{\partial T}{\partial t} = \frac{\kappa}{r} \frac{\partial}{\partial r} \left(r \frac{\partial T}{\partial r} \right) + u(t)P_v(r), \quad T(0) = T_0, \quad (1.10)$$

$$-\kappa \frac{\partial T}{\partial r} = \alpha(T - T_\infty) \quad \text{at} \quad r = R, \quad (1.11)$$

$$\frac{\partial P}{\partial t} = f(t, T), \quad P(0) = 0, \quad (1.12)$$

subject to the constraints

$$0 \leq u(t) \leq 1, \quad (1.13)$$

$$P(t^*, R) = P_{\text{desired}}. \quad (1.14)$$

1.3.2 Approaches for the control $u(t)$

In the problem description above, the control $u(t)$ can be any bounded function with values in $[0, 1]$. In order to solve the optimization problem, a discretization is necessary. This involves a discretization of the space and time, which is presented in Section 1.4, but also a discretization of the control in order to obtain an optimization problem in finite dimensions.

The following strategies to discretize the control $u(t)$ will be considered in this report:

1. A pulsed heating strategy with either equal pulse length λ for the on- and off-times or with different pulse lengths $\lambda_{0/1}$ resulting in one respectively two optimization variables.
2. Another approach is to use a smooth control $u(t)$ with parametrized decay in the form $u(t) = \exp(-\lambda t)$, where λ denotes the optimization parameter.

3. In order to allow for general heating strategies, a linear spline with m support points over the interval $[0, t^*]$ will be considered. This will result in m optimization variables.

The different results using these strategies will be presented in Section 1.5.

1.4 Numerical Solution of the Heat Equation

Optimisation algorithms usually require many evaluations of the objective function. Evaluating the objective function of the optimisation problem described in Section 1.3 requires the solution of the heat equation. Therefore, we need an efficient heat equation solver.

Space and time are discretized using finite differences. For time, we use constant time-steps, while for space we use a non-uniform grid.

1.4.1 Spatial discretisation

This non-uniformity of the spatial discretization is due to the cooling condition (1.4) imposing that the derivative of the temperature at the outer boundary being proportional to the temperature difference between the pipe and the atmosphere and the non-uniform heating pattern of the Joule heat, given by eq. (1.2). When numerically solving the heat equation with a uniform spatial grid, the temperature gradient at the boundary is fairly low, which is not representative of experimental results, but the boundary condition is physically correct so this discrepancy is put down to numerical inaccuracy. So as to not increase the cost of solving (1.3) to unreasonable levels, but to avoid the numerical inaccuracies at the boundary, a non-uniform spatial grid was used.

The spacing between discretized radii r_{i-1} and r_i , denoted by h_i is given by

$$h_i = h_{\min} + (h_{\max} - h_{\min}) \left(\frac{-(\log(i) + \log(N))}{\log(N)} + 2 \right), \quad (1.15)$$

where $h_{\min}$ is the minimum grid spacing, $h_{\max}$ is the maximum grid spacing and N is the total number of grid points. In our results, we use the following parameters: $h_{\min} = 0.01$, $h_{\max} = 1.0$, $N = 1000$. The grid is then normalized so the final spatial point is at $r = R$ by multiplying by R/r_N .

We approximate the 1st and 2nd spatial derivatives of u by

$$\begin{aligned} \left(\frac{\partial u}{\partial r} \right)_j &= \frac{1}{2(r_{j+1} - r_j)} u_{j+1} + \frac{1}{2} \left[\frac{1}{r_j - r_{j-1}} - \frac{1}{r_{j+1} - r_j} \right] u_j + \frac{1}{2} \left[\frac{-1}{r_j - r_{j-1}} \right] u_{j-1}, \\ \left(\frac{\partial^2 u}{\partial r^2} \right)_j &= \frac{2}{(r_{j+1} - r_{j-1})(r_{j+1} - r_j)} u_{j+1} + \frac{2}{(r_{j+1} - r_{j-1})} \left[\frac{-1}{r_{j+1} - r_j} + \frac{-1}{r_j - r_{j-1}} \right] u_j \end{aligned} \quad (1.16)$$

$$+ \frac{2}{(r_{j+1} - r_{j-1})(r_j - r_{j-1})} u_{j-1}, \quad (1.17)$$

In order to use the approximations (1.16) and (1.17), we use the product rule on the radial second derivative

$$\frac{\partial}{\partial r} \left(\frac{1}{r} \frac{\partial T}{\partial r} \right) = \frac{1}{r} \frac{\partial T}{\partial r} + \frac{\partial^2 T}{\partial r^2}. \quad (1.18)$$

1.4.2 Time discretization

For time-stepping, we use the Crank-Nicolson method with constant time-step Δt . It is a second order implicit one-step method, and it is numerically stable. For example, for the ordinary differential equation $y'(t) = f(t)$, and letting $y(n\Delta t) = y^n$ and $f(n\Delta t) = f^n$, the method reads as follows:

$$\frac{y^{n+1} - y^n}{\Delta t} = \frac{1}{2}(f^{n+1} + f^n). \quad (1.19)$$

In our case, at each time-step we need to solve the linear system

$$\left(I - \frac{\Delta t}{2} A \right) \mathbf{T}^{n+1} = \left(I + \frac{\Delta t}{2} A \right) \mathbf{T}^n + \frac{\Delta t}{2} (u^{n+1} + u^n) \mathbf{P}_v, \quad (1.20)$$

where $\mathbf{T}^n = [T_0^n, T_1^n, \dots, T_{N-1}^n]$ and T_i^n denotes the solution at radius r_i and time $n\Delta t$. A is the Laplacian matrix using the discretization described in Section 1.4.1.

Since the matrix $(I - \frac{\Delta t}{2} A)$ is constant, we can factorise it beforehand. In our case, we use an LU factorisation, and thus at each time-step we simply need to do forward and backward substitution.

1.5 Results

1.5.1 Pulsed Heating

The heat equation (1.3) is solved with a pulsating force term $u(t)$ with a constant pulse length λ . For each pulse, $u(t) = 1$ for a time of length λ and $u(t) = 0$ for the following λ . The solver was stopped once the desired Dorn parameter P_{desired} was reached at the outer radius R . As we are only optimising over one variable, we plot the final time at which this desired Dorn parameter is obtained over the pulse length λ , as seen in Figure 1.4. This final time can be seen to be exceedingly non-convex with many local minima. We highlight the pulse-lengths which satisfy the constraint (1.8) in green and can optimise using exhaustive search which pulse length is optimal. The end-time of the tempering is not strictly the same objective function as in (1.9) because we have neglected the penalty term and used it as a hard constraint as we only consider the green values of the end-time.




Figure 1.4: Time of heating t^* versus pulse length λ . Red plot shows the values of λ which result in a solution which does satisfy the Dorn parameter constraint. Green plot shows those that do.




Figure 1.5: Dorn gradient calculated by (1.21) against the pulse length λ .

As the final Dorn parameter at the outer boundary $r = R$ is prescribed as the finishing criterion for the solver, $P(t^*, R)$ is always known.

Due to the dimensions of the Dorn-parameter, a rescaling in the form $\tilde{P} = \log_{10}(P)$ is introduced. Then, using $\tilde{P}$ instead of P , we can reformulate the Dorn gradient constraint (1.8) as a ratio of the Dorn parameter at the outer radius $P(t^*, R)$ and that

at the hardening depth $P(t^*, R_{HD})$. Thus, we express the Dorn gradient as

$$\text{Dorn gradient} = |P(t^*, R)/P(t^*, R_{HD})|. \quad (1.21)$$

We dictate that this Dorn gradient must be less than or equal to 10, which corresponds to $\log_{10}(P(t^*, R)) - \log_{10}(P(t^*, R_{HD})) \leq 1$, i.e. we have chosen $\varepsilon = 1$ in equation (1.8). In Figure 1.5, the Dorn gradient (1.21) is plotted against the pulse length, λ . The constraint can be seen to fluctuate around the constraint level until around $\lambda = 2.0$ where any greater pulse length leads to final time solutions which violate this constraint.

Through exhaustive searching of the final time values which satisfy the Dorn gradient constraint, we find that the optimal pulse length is 1.68 which gives a heating time of 16 seconds.

Similarly, we can do pulsed heating where $u(t) = 1$ for a time length λ_1 and $u(t) = 0$ for a time length λ_0 . In this case, we have two optimisation parameters instead of one. Through exhaustive search, we obtain an optimal heating time of 12 seconds with parameters $\lambda_1 = 0.19$ and $\lambda_0 = 0.08$.

1.5.2 Exponential Heating

We now have a control of the form

$$u(t) = \exp(-\lambda t), \quad (1.22)$$

where λ is a parameter that we control. Again, we solve (1.3) numerically with this forcing term until the Dorn parameter at the surface reaches a desired point, at which point the end time t^* is measured as the quantity to be optimised. We only optimise over one parameter so we can simply optimise this by exhaustive search. In Figure 1.6, we plot the final time t^* and the Dorn gradient (1.21) against the exponent λ . We also plot the constraint on the Dorn gradient as a horizontal line at 10 and thus any final times below this line do not satisfy the constraints. As the t^* is monotonically increasing, we know that the optimal final time which satisfies the constraint must be the t^* value at which the Dorn parameter is at the constraint value. We therefore obtain an optimal exponent $\lambda = 0.03534$ which gives a heating time of $t^* = 9.95$ seconds.

1.5.3 Linear Splines

To allow for more general heating strategies, we now take $u(t)$ to be a linear spline with m support points. An example of a linear spline with $m = 11$ support points is illustrated in Figure 1.7.

Fixing a maximal time $T_{\max}$, the support points t_i are distributed with equal distances $T_{\max}/(m-1)$. Then for each of the intervals between those points, we have a parameter λ_i , $i = 0, 1, \dots, m-2$, $0 \leq \lambda_i \leq 1$, such that $u(t) = \lambda_i$ if $t_i < t \leq t_{i+1}$. The parameters λ_i are the optimization variables.




Figure 1.6: Blue plot shows the heating time taken to reach desired surface Dorn parameter. There is a maximum time imposed causing the flat line at around 0.825-0.1. Green plot shows the Dorn gradient (1.21) obtained after heating. Red plot is the level which the Dorn gradient constraint must be below.




Figure 1.7: Example of a linear spline with $m = 11$ points.

In our case, we choose $m = 101$, which gives us 100 optimization variables. To find the optimal solution, an exhaustive search is now too computationally expensive. It would be better to use a global optimization algorithm, but in the context of the study group, we use local optimization algorithms provided by the Matlab optimization toolbox. We can then try different initial guesses to find a good local minimum. The best performing optimization algorithm was SLSQP.

In Figure 1.8, we illustrate locally optimal final times t^* and the corresponding $u(t)$. We see that the best time is $t^* \approx 8.95$ for the initial guess $u(t) = 0.9$. One has to note that the problem admits local minima. Therefore, the solution of the SLSQP algorithm depends on the initial guess.




Figure 1.8: Local minima obtained with constant initial guesses.



References

- (HJ45) J. H. Hollomon and L. C. Jaffe. Time-temperature relations in tempering steel. *Trans. AIME*, 162:223–249, 1945.
- (RFG⁺10) Tamas Reti, Imre Felde, Janez Grum, Rafael Colas, Gustavo Sanchez Sarmiento, and Augusto Moita de Deus. Extension of isothermal time-temperature parameters to non-isothermal conditions: application to the simulation of rapid tempering. *Strojniski Vestnik / Journal of Mechanical Engineering*, 56(2):84 – 92, 2010.
- (RGT87) T. Réti, M. Gergely, and P. Tardy. Mathematical treatment of non-isothermal transformations. *Materials Science and Technology*, 3(5):365–371, 1987.



Problem 2. Total Hip Arthroplasty

PROBLEM PRESENTED BY: H. BRÜGGEMANN, HELSE STAVANGER

STUDY GROUP: A. AVDZHIEVA, B. LICHTBLAU, J. CHRISTMAS, JI. XU, M. GRASMAIR, T. RINGHOLM

2.1 Problem definition

In most OECD countries, the number of hip replacement operations performed has increased rapidly since 2000 ([PG14](#)) and is considered as the most effective intervention for severe osteoarthritis and hip fractures. In Sweden alone, over 17 000 hip surgeries were performed in 2016.

Total hip arthroplasty (THA) has long been used to relieve pain, improve function, and to anatomically restore normal biomechanics for patients with severe hip joint diseases, which means more precisely the center of rotation, femoral offset, and the leg length.

Before a THA operation takes place, the surgeon must plan the surgery ([AGJE15](#), [EPM98](#)) and estimate the sizes of the replacement hip components to suit the patient. In order to support planning, a standardised X-ray is taken of the pelvic area and the upper parts of the femurs. Figure [2.1a](#) shows an example of such an X-ray; figure [2.1b](#) shows the same image with key anatomical features labelled. In this X-ray it is the hip on the right-hand side of the image that is to be operated on.

In the pre-digital era, the surgeon would place a clear acetate film over the X-ray film and manually draw on it key measurement lines, a process called *templating*. X-rays are scaled representations of the truth, with the scale being determined by


the height of the emitter, and the height of the bones, above the image plate. Knowing the geometry of this setup enabled the acetate measurements to be scaled to give the true measurements the surgeons needed.

While the height of the emitter above the plate can be known precisely, different parts of the bones, such as the front and back of the pelvis, will be at different heights and therefore at different scales, and, of course, these will vary by patient. For THA the key height is that of the centre of the head of the femur. To calibrate the scaling factor, a metal sphere of known size is placed within the view of the X-ray, with its centre placed as well as possible at the same height above the plate as the centre of the femoral head. If the true diameter of this reference sphere is S (commonly 3 cm) and the measured diameter in the X-ray image is M , then all measurements made from templating that X-ray must be multiplied by a factor of S/M to obtain true sizes. The reference sphere is clearly visible towards the bottom of figure 2.1a.


Since the advent of digital X-rays, the process of templating has also been digitised, with the surgeon able to draw and move lines and prosthesis shapes over the image, and to locate and measure the reference sphere. The software then scales the measurements accordingly. Figure 2.2a shows the same X-ray as figure 2.1a with the completed digital template superimposed. The key differences between the good and bad hip joints are the distance between the femoral head and the cup that it sits in (the *acetabulum*), and the shape of the femoral head itself. The femoral head should not be in direct contact with the acetabulum, and direct contact may have caused wear on the femoral head, thereby changing its shape. Figure 2.2b shows the same patient after surgery, with the new prosthesis clearly visible. Note that no marker sphere is required in this case. For an overview of the different types of prostheses and the considerations that need to be taken into account in THA, see (VVV15).

Today, digital templating with external calibration markers has become the standard method of preoperative planning and preparing THA on digital 2D X-rays, with the aim of optimising component choice and positioning in advance of the operation, as well as minimising the risk of intra- and post-operative complications (BLR⁺15). Thus this method is a valuable process for the surgeon to get an idea of how to restore the damaged hip as well as to determine the type and appropriate size of femoral and acetabular components for the hip surgery.

The optimal position of the reference spheres is at the same vertical and lateral distance from the X-ray source compared to the hip joint centre. Clearly it is not possible to place the sphere at the optimal position, so it is usually placed either between the patient's legs or lateral to the thigh. Positioning the sphere is complicated due to the difficulty of identifying the correct anatomic landmarks by palpation, or other patient-specific factors such as obesity. In addition, in clinical routine, neither the exact height of the marker ball nor the centre of the hip is known. Overall, incorrect placement of the calibration marker may result in templating errors such as incorrect component sizes, leg lengths and offset. In general, the



(a) original X-ray




(b) annotated X-ray

Figure 2.1: Images showing (a) an original X-ray (pre-operative), and (b) the same X-ray with key features labelled.



(a) pre-operative X-ray overlaid with template



(b) post-operative X-ray showing prosthesis

Figure 2.2: Images showing (a) the original, pre-operative X-ray overlaid with templating lines, and (b) the post-operative X-ray showing the prosthesis. In (a) notice the long, almost vertical lines marking the orientations of the femurs, the short perpendicular lines to these marking the widths of the channels within the femurs, the circle marking the femoral head (and hence the centre of rotation), and the almost horizontal line joining the lower, outer corners of the obturators foramen.

accuracy of templating is often described as being to within one component size (SWJK14). Currently, there is also discussion about using 3D models for templating, e.g. with CT scans, which may be more precise and accurate compared to 2D templating with digital software or acetate templates. However, this does not seem to be a realistic alternative due to radiation exposure.

The aim of this project was two-fold: to estimate the size of the error in scaling caused by the suboptimal placing of the reference sphere, and to automate the templating process as far as possible. The remaining sections of this report are structured as follows:

2. [X-ray scaling using a reference sphere](#)

We start by quantifying the error in scale caused by the suboptimal placement of the reference sphere. An automatic process for locating and measuring the size of the sphere is described. Using the known geometry of the system, we describe a method for estimating the height of the sphere above the X-ray plate.

3. [Locating and measuring femurs](#)

Using a method related to that used for locating the reference sphere, we describe an automated process for locating the femurs and the channel that runs down the middle of them. This enables the automatic measurement of the channel width profile along the length of the femur. Part of the process of templating requires that the locations of the minor trochanters are identified; we describe the beginning of a process to automate this.

4. [Locating and measuring the femoral heads](#)

A key part of templating is the locating and measurement of the femoral heads. We describe a method for locating and measuring the good joint and then using this as a basis for doing the same thing on the damaged hip, which is more difficult. An outcome of this is the location of the centres of rotation for each hip.

5. [Identifying pelvic symmetry](#)

Part of templating is to compare/equalise each side of the body. In order to achieve this we need to establish the vertical line of symmetry through the pelvis. An automated method for doing this is described in this section.

6. [Estimating 3D shape from a 2D X-ray](#)

An X-ray is a two-dimensional (2D) projection of the 3D object. The projection causes information to be lost, so that it is impossible uniquely to reconstruct the shape of, for example, the pelvis from the X-ray. Such a reconstruction could potentially enable surgeons to design unique, 3D-printed components for THA. In this section we use a "standard" hip as the basis for the 3D reconstruction and morph it with information from the 2D X-rays.

Conclusions are drawn in section [2.7](#).

2.2 X-ray scaling using a reference sphere

In this section we consider the placement and automatic location and measurement of the reference sphere. The reference sphere is most likely to be in a suboptimal position laterally, as the ideal position is occupied by the patient. We start by describing an automatic process for locating and measuring the size of the sphere. We then quantify the scaling error that results from this suboptimal placement.

A number of researchers have investigated the errors associated with using a reference sphere for scaling calibration (e.g. (ACT⁺16, BLR⁺15, FGH10)). Archibeck et al. (ATC⁺17) go so far as to suggest that a standard scaling of 21% is applied in all cases. The location of the reference sphere laterally on the X-ray plate has an effect on accuracy (RFE⁺18, SWJK14), but we show that the effect is small.

Using the known geometry of the system, we also describe a method for estimating the height of the sphere above the X-ray plate. This information could be known *a priori*, as it can be measured by the radiologist at the time the sphere is placed, but it has not generally been recorded. While this information may not be directly useful for templating, it may provide interesting statistical information when compared with, for example, patients' body mass index (BMI). In particular, this could be used in order to detect irregularities during the recording of the X-ray image.

2.2.1 Automatic locating and sizing of the sphere

The Circular Hough Transform (CHT) (Bal87, DH72, Hou61) aims to find circles of a given radius within an image; to find circles of unknown radii, the traditional CHT is performed for each possible radius. A particular variant of CHT called Phase Encoding (AK99) has been used in this report; it locates both the centres and radii simultaneously. Figure 2.3 shows the automatically detected reference sphere in four example X-rays.

2.2.2 Scale error caused by sphere location

In the software that is currently used, it is assumed that the reference sphere appears as a solid circle of the same radius in the X-ray image. Thus it can be used for converting the pixel scale of the X-ray image to a centimetre scale. A closer inspection of the geometry of this set-up, however, shows that this only holds (approximately), if the ball is placed in the centre of the X-ray image. Otherwise it will appear as an ellipse with eccentricity depending on its distance from the centre and also the distance from the X-ray source. Because of that, even a placement of the marker sphere at the correct height can lead to an error in the scale estimation, if the sphere is placed near the edge of the X-ray image. Figure 2.4 shows a sketch of the geometry of that situation; note that the sketch is not to scale.




Figure 2.3: The automatically detected reference sphere, marked in magenta, in four example X-ray images.

Using the notations in Figure 2.4 we use the following notation:

$$\begin{aligned}
 SA &= y, & AP &= t, & PD &= PC = r, \\
 AK &= s_1, & AB &= s_2, \\
 \tan(\theta + \phi) &= \frac{t}{y}, & SP &= \sqrt{t^2 + y^2}, & \sin \theta &= \frac{r}{SP}.
 \end{aligned}$$

Then the marker sphere (of diameter $2r$) has the same appearance on the X-ray image as an ellipse of maximal diameter $s_2 - s_1$ that is situated at the same height as the marker sphere. The relative scaling error ε is therefore given by the ratio

$$\varepsilon = \frac{s_2 - s_1 - 2r}{2r}.$$

$$- \tan \left(\arctan \frac{t}{y} - \arcsin \frac{r}{\sqrt{t^2 + y^2}} \right) \quad (2.8)$$

If the ratio between the radius r and the vertical distance y between the X-ray source and the marker sphere is small (in practice, this ratio is about 0.02), then we can approximate this error by

$$\varepsilon \approx \frac{\sqrt{y^2 + t^2} - y}{y}.$$

A standard X-ray plate is a square with sides of 40 cm, and the sphere is placed about 40 cm below the X-ray source. Figure 2.5 shows a plot of the percentage error underestimate of the size of an object in the sphere's plane by distance of the sphere from the centre of the X-ray plate. When the sphere is placed at the very edge of the plate, the size of the objects is underestimated by about 1.65%, which is probably not significant.




Figure 2.5: A plot showing how the percentage error underestimate of the size of an object in the sphere's plane by the lateral distance of the sphere from the centre of the 40×40 cm X-ray plate.

2.2.3 Estimation of the elevation above the X-ray plate

Figure 2.4 shows the geometry of the reference sphere projected onto the X-ray plate. From this, and given the known height of the X-ray emitter above the plate (the distance marked by the line OS in the diagram) and the size of the sphere, we may estimate the height of the reference sphere above the plate. A similar approach is described by Boese et al. (BBD⁺16).

We define the following distances:

$$OS = h \quad (2.9)$$

$$AO = z \quad (2.10)$$

$$SC = w \quad (2.11)$$

$$CB = v \quad (2.12)$$

Working through the geometry we get:

$$\frac{w}{r} = \cot(\theta) \quad \Rightarrow \quad w = r \cot(\theta) \quad (2.13)$$

$$\frac{v}{r} = \tan(\phi + 2\theta) \quad \Rightarrow \quad v = r \tan(\phi + 2\theta) \quad (2.14)$$

$$\frac{h-z}{w+v} = \cos(\phi + 2\theta) \quad \Rightarrow \quad z = h - (w+v) \cos(\phi + 2\theta) \quad (2.15)$$

which gives us an expression for z based on the angles ϕ and θ :

$$z = h - r \cos(\phi + 2\theta) \left(\cot(\theta) + \tan(\phi + 2\theta) \right) \quad (2.16)$$

With these angles defined as

$$\phi = \arctan\left(\frac{x_1}{h}\right) \quad (2.17)$$

$$\theta = \frac{1}{2} \left[\arctan\left(\frac{x_2}{h}\right) - \arctan\left(\frac{x_1}{h}\right) \right] \quad (2.18)$$

we end up with the following expression for z in terms of known distances and the distances we can measure:

$$z = h - r \cos\left(\arctan\left(\frac{x_1}{h}\right)\right) \left(\cot(\theta) + \frac{x_2}{h} \right) \quad (2.19)$$

We performed three experiments to find the height of the reference sphere in two real images with different assumed heights of the emitter. In general the emitter height is known, but we do not have the information for these X-rays. The results are shown in table 2.1 below, where, the elevation of the reference sphere is denoted by z and it is a function of the values of the radius r , the height h , the angle θ , and distances x_1 and x_2 .

We know that the surface on which the patient lies during the X-ray is 7 cm above the X-ray plate (Brü18), so the distance from that surface to the centre of the sphere, and hence, presumably, to the centre of the heads of the femurs is estimated to be of the order of 15 cm, which does not seem unlikely.

It will be interesting to plot this distance against BMI for a large number of patients and to derive a statistical model from the data. It would then be possible to check the likelihood of the distance estimate for a given patient X-ray, given their BMI and the statistics from the cohort, and thus possibly highlight cases where the sphere has been placed at the wrong height for that patient.

image	emitter height h	sphere radius r	distance x_1	distance x_2	estimated sphere height z
1	110	1.5	16.38	20.30	24.64
1	100	1.5	16.38	20.30	22.18
2	110	1.5	4.89	8.73	23.88
2	100	1.5	4.89	8.73	21.68

Table 2.1: Estimation of the sphere’s elevation, z , above the X-ray plate, given the key distances defined in the geometry shown in figure 2.4. All distances are in centimetres.

2.3 Locating and measuring femurs

The hip implant is typically anchored in the patient’s femur. In order to partially automate the templating process, one can obtain the location and orientation of the femurs and use this information to (a) align the orientation of the implant with the orientation of the femur, and (b) obtain a width profile of the femur to help choose an implant of the right shape and size. An accurate measurement of the femur channel profile is particularly important for cementless implants (NLM18).

In this section we demonstrate the automatic location of the femurs and their orientation, from which measurements of the width profile of the channel that runs down their centres may be made. A key part of the templating process is to locate the two minor trochanters towards the top of the inner edge of the femurs; we demonstrate a method for achieving that.

2.3.1 Automatic location and orientation

The Hough transform (DH72, Hou61) can be used for the detection of lines within an image. By first applying an edge detection filter and then a line detection step using MATLAB’s `hough`, `houghpeaks` and `houghlines` functions with properly chosen parameters, we can identify the outside edges of the femur, from which we can estimate middle lines as shown in figure 2.6. Note that this figure only shows a proof of concept for a particular image with a particular set of parameters used in the MATLAB functions, and that the process as of now does not automate to all similar images.

2.3.2 Estimation of channel widths

Once the femoral midlines are identified, we can find estimates of the channel widths by considering the distances between the inside edges of the bones measured by lines perpendicular to the femoral midlines, as presented in figure 2.7. The rightmost panel in figure 2.7 shows the channel width estimated in pixel numbers, which can be converted to metric units using a conversion after proper calibration.




Figure 2.6: Femurs with outer edges marked in green and central orientation line in yellow.




Figure 2.7: Left: Rotation of left hand femur. Middle: Cutout of a part of the left hand femur. Right: Channel width, in pixels, of the left femur.

2.3.3 Automatic location of minor trochanters

It would be nice if we could just use standard edge detection techniques to find the outline of the bone mass, as a starting point for identifying the minor trochanters and hence their vertical offset. Figure 2.8 shows the original X-ray and the result of three different edge detection algorithms on that X-ray. We can see that although parts of the edges have been identified (the Sobel edge detector in (b) has given the clearest results), there is no good definition of the bone mass. In all three cases detectors have also identified non-bone edges, and, particularly for the Canny and Log edge detectors, the resulting images are very noisy.

There are many different algorithms available (OH10, Bas02), and they all require a certain amount of manual tuning for a particular application, and even for different instances of the same application (for example, X-rays taken by different machines).

In the field of computer vision, *contour models* aim to learn the locations of the boundaries of shapes in an image. As a first step towards locating the minor trochanters, we used an active contour model called *Snakes* (KWT88) to try and find the edges of the bones. The algorithm progressively moves and deforms a spline until it reaches some sort of boundary where the “energy” required to overcome the boundary is too high. Snakes has been applied to THA templating, but apparently only to locate the reference sphere (KvOD⁺07).

An initial test of this algorithm against the same X-ray image is shown in figure 2.9a. The Snakes spline was initialised along the edges of the image and it then moved progressively inwards, towards the centre of the image, until it encountered the edges of the bones. The resulting outline is very smooth, in fact too smooth; the minor trochanters have been smoothed away because their edges are not well-defined enough to present a boundary to the Snakes algorithm. To get around this, we augmented the original X-ray image with the detected edges shown in figure 2.8b, by simply adding them together to obtain the image in figure 2.9b, and then performed the Snakes algorithm against that augmented image. The result is shown in figure 2.9c. It has worked well on the femurs, and the minor trochanters are now visible.

In this example, the Chan-Vese (CV01) version of the Snakes algorithm was used, with 20 000 iterations and a smoothing factor of 2.

Starting from the bottom of the bone mass image in figure 2.9c, we can work upwards until a row of pixels is reached that contains more than two bones. This enables us to extract the two femurs (although not the minor trochanters) and make some measurements of their diameters. Figure 2.10 shows the extracted femurs and plots of the femur width along the lengths of the bones. Note that this calculation is not exactly correct as it traces the width across each row of pixels, whereas it should calculate the width across the bone, perpendicular to the femoral midlines.

These are only the first steps towards locating the minor trochanters. We have shown that the edges of the femurs can be located using this method. Following




Figure 2.8: Figure (a) shows the original X-ray; the other three images show the results of three different edge detection algorithms performed on that X-ray. Note that the marker ball was automatically detected and removed from the image as it presents a very clear edge for all the algorithms described in this section.

the smooth inner edges of the bones upwards, the minor trochanters represent a significant deviation away from this smoothness, so the next step should be straightforward.




Figure 2.9: Figure (a) shows the outline of the bone mass identified by the Snakes algorithm on the original X-ray. The minor trochanters have been smoothed away because their edges are not well-defined enough to present a boundary to halt the algorithm. Figure (b) shows the original X-ray from figure 2.8a augmented with the edges detected by the Sobel edge detector, shown in figure 2.8b. Figure (c) shows the outline of the bone mass identified by the Snakes algorithm on this augmented image. Now the minor trochanters are clearly visible.




Figure 2.10: The image on the left shows the femurs extracted from 2.8a as a result of the identification of the bone mass using the Snakes algorithm. The graph on the right shows the extracted bone diameters for the length of each femur.

2.4 Locating and measuring the femoral heads

In a healthy hip joint, the top of the femoral head has a perfectly circular shape and is separated from the acetabulum, in which it is embedded, by a small layer of cartilage. As a consequence, it essentially appears in the X-ray image as an

almost solid, bright circle (the femoral head proper consisting of bone), separated from another bright semi-circle (the acetabulum, also consisting of bone) by a dark region (the cartilage in between, with significantly lower X-ray absorption). The thickness of the cartilage layer, and the size and center of the femoral head can vary, but the general structure is necessarily the same in a functioning hip. As a consequence, it should be possible to find the center and size of the femoral head by searching for parts of the X-ray image that exhibit the same set of features.

To that end, we define for some possible center position (x, y) of the femoral head, and radii $r_1 < r_2 < r_3 < r_4$ the regions (see Figure 2.11)

$$\begin{aligned}\Omega_0 &= (x, y) + \{t(\sin \vartheta, \cos \vartheta) : -\pi/4 \leq \vartheta \leq \pi/4, t \in [r_1, r_2] \cup [r_3, r_4]\}, \\ \Omega_1 &= (x, y) + \{t(\sin \vartheta, \cos \vartheta) : -\pi/4 \leq \vartheta \leq \pi/4, t \in [r_2, r_3]\}.\end{aligned}$$

Moreover, we denote by I_0 and I_1 the mean intensities of the X-ray image in the regions Ω_0 and Ω_1 , respectively. Then we vary the center (x, y) and the radii $r_1 < r_2 < r_3 < r_4$ until the difference $I_0 - I_1$ becomes maximal, that is, until we obtain the largest possible contrast between these regions. The point (x, y) should then be a good approximation of the center of rotation, and the radii r_2 and r_3 are approximately equal to the radius of the femoral head (in pixels) and the radius of the cup of the acetabulum.




Figure 2.11: Sketch of the two masks used for the detection of the femoral head.

In order to make this approach work, it is necessary to restrict the possible radii r_i and the centers of rotation to some reasonable search space. Concerning the radii, the numerical experiments showed good results for a constant choice of the differences $r_2 - r_1 = r_3 - r_2 = r_4 - r_3 = 3$ pixels (for image sizes of roughly 800×800 pixels — the values have to be adjusted to the resolution of the image). Thus, one only needs to maximise the contrast difference over the parameter r_2 describing the radius of the femoral head. Here, any reasonable restrictions of the size of the femoral head gave consistent results. However, it is necessary to restrict the search space for the center position (x, y) , as there can be different features in the image that either have a similar structure, or are circular with stronger contrast; the latter holds in particular for the marker ball. If the search space for the center is roughly restricted to the region of the femoral head, though, the position of the center can be found reliably. In the current implementation, bounds for the search space have to be entered manually; however, it seems reasonable that an automatic selection




Figure 2.12: Semi-automatic detection of center of rotation. In a first step, the center of rotation and radius of the femoral head is detected on the healthy side. Then the estimated size of the femoral head is used in order to find the center of rotation on the damaged side. The rough position of the femoral heads has to be provided by the user.

of the region of interest can be implemented that uses information on the femur positions and orientations obtained in the previous step.

In case of the damaged side, the situation is more complicated, as the gap between femoral head and acetabulum is either much smaller or non-existent. Still, experimental results indicate that the same contrast maximisation procedure is able to find the position of the center of rotation, provided it is given a reasonable estimate of the position of that center as well as an accurate estimate of the radius of the femoral head. The latter can be obtained by using the estimate from the healthy side; the former is in the current implementation chosen manually, but we expect again that an automatic selection is possible that uses the results of the previous computations.

Examples of the results of the contrast maximisation are depicted in figure 2.12.

2.5 Identifying pelvic symmetry

Many of the measurements around the alignment of the hips and orientations of the femurs are dependent on establishing a vertical centre-line to identify the line of symmetry of the pelvis. In general all the X-rays are oriented in approximately the same way, with the pelvis at the top and femurs at the bottom, but the condition that leads to the patient being considered for a hip replacement can also lead them to have difficulty lying symmetrically on the X-ray table. For example, they may not be able to rotate both hip joints outwards. In this section we propose a method for automatically determining a vertical axis of symmetry for each X-ray.

Speeded Up Robust Features (SURF) ([BETV08](#)) is an algorithm that finds “interesting”

features in an image, features such as corners, edges and blobs. For each of these features it extracts a characteristic *feature vector* that describes it and its immediate neighbourhood. If we have two similar photographs of the same object, the algorithm can match pairs of feature vectors (one from each image) that are the most similar. The set of matches as a whole can be used to identify how an object's orientation has changed between the two images. SURF is a very quick algorithm, able to match two X-ray images in less than half a second on a standard laptop.

The algorithm developed during the workshop proceeds as follows, with the letters in brackets corresponding to the subfigure in figure 2.13 that shows an example output from the step:

- (a) start with the original X-ray image
- (b) split the image in half, vertically
- (c) flip the right-hand half horizontally so that both halves of the image now look similar
- (d) for each half, use SURF to identify the interesting features and extract feature vectors
- (e) match pairs of feature vectors from each half
- (f) join each pair of matched features with a line and find their centres
- (g) fit a straight line to these centre points
- (h) this centre-line is the axis of symmetry

In this example features from the femurs are included, which is probably not appropriate. They can be excluded from the process by using results section 2.3.3.

Figure 2.14 shows the resulting vertical axis of symmetry for five different X-rays. The sixth example has been artificially rotated to show that the algorithm works even when the orientation is quite markedly offset from the centre-line of the X-ray.

2.6 Estimating 3D shape from a 2D X-ray

An X-ray is a two-dimensional (2D) projection of the 3D object. The projection causes information to be lost, so that it is impossible uniquely to reconstruct the shape of, for example, the pelvis from the X-ray. Such a reconstruction could potentially enable surgeons to design unique, 3D-printed components for THA, and there is evidence that 3D templating can lead to better pre-operative planning outcomes (MBK⁺17). In this section we use a "standard" hip as the basis for the 3D reconstruction and morph it with information from the 2D X-rays.

The section is split into three parts. First we define how X-ray makes a 2D projection of the pelvis. Then we consider the inverse of this process: making a 3D reconstruction from a 2D projection. Finally we look at morphing a standard pelvis shape to fit an individual patient's X-ray to potentially enable 3D templating from the 2D X-rays.




Figure 2.13: The results of each step of the SURF feature matching process for a single X-ray, showing how the vertical axis of symmetry is estimated.




Figure 2.14: The vertical axis of symmetry estimated for five different X-rays. The bottom right image is the same as the bottom left, but has been artificially rotated before the matching process to show that the process still works when the axis of symmetry is markedly offset from the centre-line of the X-ray.




Figure 2.15: 2D projection of 3D pelvis model

2.6.1 3D to 2D projection

The process of taking an X-ray image can be simulated as taking a perspective projection of a 3D object onto a 2D plane. Here the vanishing point is at the X-ray source position, the 2D plane is the radiograph, and the aspect ratio and the field of view angle are determined by the size of radiograph and its distance to the X-ray source, see figure 2.15.

With a 3D model of pelvis and its position and orientation in space known to the observer. Its 2D radiograph can be simulated under different configurations of X-ray source and radiograph.

The information a 3D pelvis model contains includes but is not limited to vertex position, normal vectors, surface indices, and voxelized material properties.

Denote the vertex positions in model reference frame as

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix}$$

which is constant regardless of the overall position and orientation of the 3D model

The projected vertex position on 2D plane as

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix}$$

Here z' represents the depth information of a 2D projection.

The position and orientation of the model can be represented as a 4-by-4 transformation matrix known as the “model matrix”

$$M_{model} = \begin{bmatrix} R_{xx} & R_{xy} & R_{xz} & P_x \\ R_{yx} & R_{yy} & R_{yz} & P_y \\ R_{zx} & R_{zy} & R_{zz} & P_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

where R represents 3-by-3 rotation matrix of orientation, and P represents 3-by-1 translation in world space.

The second matrix we need is the “view matrix”, which represents the position and viewing direction of the X-ray source.

$$M_{view} = \begin{bmatrix} R_x & R_y & R_z & -Q_x \\ U_x & U_y & U_z & -Q_y \\ D_x & D_y & D_z & -Q_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

where R is the right vector, U is the up vector, D is the direction vector, and Q is the X-ray source position. If we set the origin of reference frame at x-ray source and $-z$ axis pointing towards the radiograph, the view matrix can be simplified as an identity matrix, which can be omitted.

The third matrix is the “projection matrix”, which represents the perspective projection effect from 3D to 2D.

$$M_{projection} = \begin{bmatrix} \frac{1}{aspect \cdot \tan(0.5 \cdot fov)} & 0 & 0 & 0 \\ 0 & \frac{1}{\tan(0.5 \cdot fov)} & 0 & 0 \\ 0 & 0 & -\frac{far+near}{far-near} & -\frac{2 \cdot far \cdot near}{far-near} \\ 0 & 0 & -1 & 0 \end{bmatrix}$$

where

$$aspect = \frac{width}{length}$$

of the radiograph, fov is the field of view angle

$$\tan(0.5 \cdot fov) = \frac{length}{2 \cdot height}$$

The *height* is the distance from radiograph to the x-ray source. *far* and *near* are two control parameters for the depth data.

The final equation for 3D to 2D x-ray projection is

$$M_{projection} \cdot M_{view} \cdot M_{model} \cdot \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix}$$

2.6.2 2D to 3D reconstruction

In reality we only have the information of the 2D radiograph and the machine configuration. A 2D image to 3D model reconstruction can be conducted with the extra information from the self-symmetry of the pelvis and previously derived marker ball position in 3D, see figure 2.16.

Given a 2D radiograph, some useful relationship can be established, see figure 2.17. Let us assume that on the radiograph, two non-collinear pairs of symmetrical reference points have been identified by 2D image processing(symmetry detection), denoted as point A, B and C, D . Line AB and line CD are not necessarily parallel to each other, but their actual position $A'B'$ and $C'D'$ in 3D space before projection should be parallel because of the symmetry of the pelvis. This is the first symmetry.

$$A'B' \parallel C'D'$$

The second symmetry is that the distance between two symmetrical points on each side of the pelvis should be identical

$$|A'D'| = |B'C'|$$

Finally, all 3D points A', B', C', D' are located on the rays from x-ray source O to their projected points on 2D radiograph.

Using aforementioned geometrical relationships, we can establish several equations. Set the plane function of $A'B'C'D'$ as

$$\alpha x + \beta y + z + \delta = 0$$

The relationships of the normal vector of each plane are

$$\vec{N}_{OAB} = \vec{OA} \times \vec{OB} \quad (2.20)$$

$$\vec{N}_{OCD} = \vec{OC} \times \vec{OD} \quad (2.21)$$

$$\vec{\Phi} = [\alpha; \beta; 1] \quad (2.22)$$

$$\vec{N}_{A'B'} = \vec{N}_{OAB} \times \vec{\Phi} \quad (2.23)$$

$$\vec{N}_{C'D'} = \vec{N}_{OCD} \times \vec{\Phi} \quad (2.24)$$

$$A'B' \parallel C'D' : \vec{N}_{A'B'} \times \vec{N}_{C'D'} = 0 \quad (2.25)$$

$$\vec{\Phi} \cdot (\vec{N}_{OAB} \times \vec{N}_{OCD}) = 0 \quad (2.26)$$




Figure 2.16: 3D symmetrical pelvis with marker ball




Figure 2.17: 3D reconstruction from 2D radiograph

$$|A'B'| = |C'D'| \quad (2.27)$$

Equations 2.26 and 2.27 only contain unknown variables α and β , meaning we can determine the orientation of the 3D pelvis model without the help of the marker ball. The use of marker ball is to determine the scale of the 3D model. The prerequisite is that the marker ball has to be located at the same height as one of the symmetrical point. For example,

$$z_{\text{marker}} = z_{C'}$$

Now δ can also be found.

With more pairs of symmetrical points identified by 2D image processing, more 3D plane can be found by the aforementioned approach, eventually creating a set of anchor points for a standard 3D model.

2.6.3 2D/3D hybrid templating

Assuming that the patient has a symmetrical pelvis, we can fit a standard 3D pelvis model to the anchor points found in 2D to 3D reconstruction, and conduct 2D/3D hybrid templating. The procedure can be described as follows:

1. Position the marker ball at the same elevation as one of the key reference points.
2. Find the symmetrical pair of that key reference point on the 2D radiograph.
3. Find other pairs of symmetrical reference points on the 2D radiograph.
4. Find the elevation of the marker ball
5. Use the information of symmetrical pairs, marker ball elevation, radiograph, and X-ray source to reconstruct 3D coordinates of those symmetrical pairs.
6. Use reconstructed 3D coordinates as anchor points to fit in a standard 3D pelvis.
7. Simulate the perspective projection of the fitted 3D pelvis.
8. Compare the original 2D radiograph and the simulated radiograph.
9. Micro adjust the 3D pelvis for a good overlapping fit.
10. Conduct templating on 3D model and calibrate with simulated 2D radiograph

2.7 Conclusions and Future Work

In this work we have shown that the lateral position of reference sphere produces probably insignificant errors in the calibration of the X-ray scale compared with when it is placed in the optimal position.

We have also shown a variety of methods for automating different parts of the digital templating process. All of these depend on a certain amount of tuning in order to find parameters that will suit the standard pre-operative THA X-rays, though it is likely that there will always be individual situations where the automation will fail. However, even just using the automation to make a first estimate of the template, which the

surgeons can use their expertise to adjust appropriately, would save the surgeons' time at a key moment in the procedure. There is evidence that the expertise of the surgeon in the templating process is key to its success ([SPR17](#), [MIS⁺12](#)), a view supported by ([PSO⁺15](#)), who find that using manual, acetate templating over scaled digital X-rays is more accurate than digital templating.

Acknowledgements

Ana Avdzhieva acknowledges for the support provided by the Sofia University Research Fund through Contract No. 80-10-139/2018.



References

- (ACT⁺16) Michael J Archibeck, Tamara Cummins, Krishna R Tripuraneni, Joshua T Carothers, Cristina Murray-Krezan, Mohammad Hattab, and Richard E White. Inaccuracies in the use of magnification markers in digital hip radiographs. *Clinical Orthopaedics and Related Research*, 474(8):1812–1817, 2016.
- (AGJE15) Kivanc Atesok, David Galos, Laith M Jazrawi, and Kenneth A Egol. Preoperative planning in orthopaedic surgery. *Bulletin of the Hospital for Joint Diseases*, 73(4), 2015.
- (AK99) T.J. Atherton and D.J. Kerbyson. Size invariant circle detection. *Image and Vision Computing*, 17:795–803, 1999.
- (ATC⁺17) Michael J Archibeck, Krishna R Tripuraneni, Joshua T Carothers, Daniel W Junick, Natalie R Munson, and Cristina M Murray-Krezan. Prospective, randomized, surgeon-blinded comparison of standard magnification assumption vs magnification marker usage for preoperative templating in total hip arthroplasty. *The Journal of Arthroplasty*, 32(10):3061–3064, 2017.
- (Bal87) Dana H Ballard. Generalizing the hough transform to detect arbitrary shapes. In *Readings in computer vision*, pages 714–725. Elsevier, 1987.
- (Bas02) M. Basu. Gaussian-based edge-detection methods – a survey. *IEEE Transactions on Systems Man and Cybernetics*, 32(3):252–260, 2002.
- (BBD⁺16) Christoph K Boese, Jan Bredow, Jens Dargel, Peer Eysel, Hansjörg Geiges, and Philipp Lechler. Calibration marker position in digital templating of total hip arthroplasty. *The Journal of Arthroplasty*, 31(4):883–887, 2016.

- (BETV08) H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool. Speeded-Up Robust Features (SURF). *Computer Vision and Image Understanding*, 110(3):346–359, 2008.
- (BLR⁺15) Christoph Kolja Boese, Philipp Lechler, Leonard Rose, Jens Dargel, Johannes Oppermann, Peer Eysel, Hansjörg Geiges, and Jan Bredow. Calibration markers for digital templating in total hip arthroplasty. *PloS one*, 10(7):e0128529, 2015.
- (Brü18) H. Brüggemann. personal email, June 2018.
- (Can86) J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.
- (CV01) T.F. Chan and L.A. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.
- (DH72) Richard O Duda and Peter E Hart. Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1972.
- (EPM98) S Eggli, M Pisan, and ME Müller. The value of preoperative planning for total hip arthroplasty. *Journal of Bone & Joint Surgery (Br)*, 80(3):382–390, 1998.
- (FGH10) M Franken, B Grimm, and I Heyligers. A comparison of four systems for calibration when templating for total hip replacement with digital radiography. *Journal of Bone & Joint Surgery (Br)*, 92(1):136–141, 2010.
- (Hou61) P.V.C. Hough. Method and means for recognizing complex patterns, 1961.
- (KvOD⁺07) W Kristanto, Peter MA van Ooijen, H Duifhuis, TR Mengko, M Oudkerk, et al. Design and implementation of gradient vector flow snake to detect a reference object in pelvic X-rays for preoperative total hip arthroplasty planning application. *Journal of Digital Imaging*, 20(4):373–380, 2007.
- (KWT88) M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.
- (MBK⁺17) D Mainard, O Barbier, Y Knafo, R Belleville, L Mainard-Simard, and J-B Gross. Accuracy and reproducibility of preoperative three-dimensional planning for total hip arthroplasty using biplanar low-dose radiographs: A pilot study. *Orthopaedics & Traumatology: Surgery & Research*, 103(4):531–536, 2017.
- (MH80) D. Marr and E. Hildreth. Theory of edge detection. *Proceedings of the Royal Society of London Series B*, 207:187–217, 1980.

- (MIS⁺12) Falk Mittag, Ingmar Ipach, Ralf Schaefer, Christoph Meisner, and Ulf Leichtle. Predictive value of preoperative digital templating in tha depends on the surgical experience of the performing physician. *Orthopedics*, 35(2):e144–e147, 2012.
- (NLM18) Shahryar Noordin, Riaz Lakdawala, and Bassam A Masri. Primary total hip arthroplasty: Staying out of trouble intraoperatively. *Annals of Medicine and Surgery*, 2018.
- (OH10) M.A. Oskoei and H. Hu. A survey on edge detection methods. techreport CES-506, University of Essex, 2010.
- (PG14) C Pabinger and A Geissler. Utilization rates of hip arthroplasty in oecd countries. *Osteoarthritis and cartilage*, 22(6):734–741, 2014.
- (PSO⁺15) Robert Petretta, Jason Strelzow, Nicholas E Ohly, Peter Misur, and Bassam A Masri. Acetate templating on digital images is more accurate than computer-based templating for total hip arthroplasty. *Clinical Orthopaedics and Related Research*, 473(12):3752–3759, 2015.
- (RFE⁺18) Austin J Ramme, Nina D Fisher, Jonathan Egol, Gregory Chang, and Jonathan M Vigdorchik. Scaling marker position determines the accuracy of digital templating for total hip arthroplasty. *HSS Journal*, 14(1):55–59, 2018.
- (Sob14) I. Sobel. History and definition of the so-called “Sobel Operator”, more appropriately named the Sobel-Feldman Operator. https://www.researchgate.net/publication/239398674_An_Isotropic_3x3_Image_Gradient_Operator, last accessed 14 October 2018, February 2014.
- (SPR17) Nils J Strøm, Are Hugo Pripp, and Olav Reikerås. Templating in uncemented total hip arthroplasty—on intra-and interobserver reliability and professional experience. *Annals of Translational Medicine*, 5(3), 2017.
- (SWJK14) Victoria F Sinclair, James Wilson, Neil PM Jain, and David Knowles. Assessment of accuracy of marker ball placement in pre-operative templating for total hip arthroplasty. *The Journal of Arthroplasty*, 29(8):1658–1660, 2014.
- (VVVV15) Jan Vanrusselt, Milan Vansevenant, Geert Vanderschueren, and Filip Vanhoenacker. Postoperative radiograph of the hip arthroplasty: what the radiologist should know. *Insights into Imaging*, 6(6):591–600, 2015.



Problem 3. Predicting salmon weight

PROBLEM PRESENTED BY: ARNE GUTTVIK, SALMAR

STUDY GROUP: H. BATARFI, Y. CAO, P.G. HJORTH, D. HÖMBERG, S. VAN ALBADA, J. TYSSDAL

3.1 Introduction

The company SALMAR is one of Norway's largest salmon producers, contributing significantly to Norway's export and to the global ocean-farm-produced salmon

In the production, the salmon are kept in large circular cages ('pens') in the ocean, and fed over the course of 18-24 months, after which they are slaughtered and sold ([MR01](#)). It is important to be able to predict the distribution of weights for the total number of salmon in the pen. This distribution is rarely normal, and in fact is often seen to develop bi-modality ([JSW85](#)).

As the population increases in weight, samples are taken in order to check the evolving distribution. Since the sample is finite (and the distribution is finite), a number of fluctuations are present and the predicted weight distribution comes with a considerable uncertainty([FAA+16](#)). The questions put to the Study Group were:

- Is it possible to deal with uncertainties in a quantitative way?
- How could we build a mathematical model that will increase the accuracy of prediction for the weight distribution?
- Is it possible to combine information from weight samples and amount of feeding to give a better prediction of weights or early warning of abnormal

Weight density functions




Figure 3.1: Weight distributions (density functions) for different pens. The distributions are far from normal, and often develop bi-modality

distribution developing?

3.2 Deterministic analysis

The classical ODE model for the growth of an isolated population (or, equivalently, growth of the average weight of an isolated population) is based on the logistic equation. The average weight $w(t)$ of a salmon population as a function of time, beginning with the initial value w_0 , is taken as the solution to a logistic equation

$$\frac{dw}{dt} = K\left(1 - \frac{w}{k}\right)w, \quad w(0) = w_0 \quad (3.1)$$

The solution curve for this equation is the ideal, or continuous, growth curve, given K and k . The two parameters K and k signify the initial growth rate, and the asymptotic (i.e., $t \rightarrow \infty$) weight respectively, for the population. Sometimes k is known as the *capacity* for the population. On figure 3.2 one can see the characteristic s-growth of the weight with time, for various initial conditions.




Figure 3.2: Ideal growth curves (average weight as a function of time) as solutions of the logistic equation (3.1) for a various initial values. Here, the capacity k has been normalised to 1.

Both of the parameters K and k are for a given population determined by a number of factors that will vary from season to season, and are therefore only approximately known at the outset for a given population.




Figure 3.3: Growth curves of the logistic equation (3.1) with different values of the growth parameter K . Only at the later stages will uncertain measurements be able to distinguish between the curves.

Note that two solution curves with the same k but different K 's can look very similar at the outset (see figure 3.3). This means that K is poorly determined by measurements on the early part of the weight evolution curve.

3.3 Statistical analysis

3.3.1 Preliminary analysis

Figure 3.4 is scatter plot of individual weight against length of salmons from two farming cages. It shows strong correlation between these two phenotypic traits (correlation=0.98). As a consequence there may not be much point in measuring both length and weight. For slaughtered fish the estimated coefficient of variation, CV, was given for about 600 cages and was found to be approximately constant as a function of the average weight. Though for a given weight there was large variation. This implies that the standard deviation given as $\sigma = CV\mu$ can be assumed to be proportional with μ or for practical purposes the average weight. The estimate for CV or the proportional factor was about 0.22.

3.3.2 Sampling consequences

Constant CV and constant sample size independent of the weight of the fish implies that % standard deviation is constant, but the uncertainty increases with the weight. This again implies that the information from the samples becomes smaller when the weight of the fish increase. This is a problem for the prediction of the weights at slaughter. The sampling procedure is also more frequent in the beginning than at the end where usually only two samples are taken. The assumption of a growth curve following the sigmoid curve, may be a reason for the frequent sampling in

the beginning. Though it is hard to predict where the fish-population is on this curve when slaughtered, and the possibility that the fish grow into two or more populations due to sexes and diseases is an argument for more sampling at the end.

3.3.3 The effect of having two populations

Expectedly weights in populations are reasonable well approximated by a normal distribution. If two populations are present in the cage, a mixture of two normal distributions seems to be an appropriate model for the weights. Assume $X_1 \sim N(\mu_1, \sigma_1^2)$ and $X_2 \sim N(\mu_2, \sigma_2^2)$ and that p is the proportion of the first population. Then for an arbitrary fish, Y , its expected weight and variance is given by:

$$E(Y) = p\mu_1 + (1 - p)\mu_2$$

$$Var(Y) = p\sigma_1^2 + (1 - p)\sigma_2^2 + p(1 - p)(\mu_1 - \mu_2)^2.$$

The differences in the means will normally affect the variance the most. An illustration of the mixture distribution is demonstrated by the simulator in Section 3.5.




Figure 3.4: A plot of individual salmon weight (g) against length.

3.3.4 Statistical model

We constructed a statistical model to forecast, with uncertainty, future (especially when slaughtering) biomass of farmed salmon.

Based on the available data, we fitted a linear regression model with historical salmon data from 24 farming cages. For each cage, we have the mean weight of fish at slaughtering, the mean weight of samples (the samples were not drawn at fixed time with fixed number of fish, the time ranges from one month to approximate three months before slaughtering and sample size was typically 60 or 100). Other variables include the sea temperature when slaughtering and the weight difference (we have data of predicted mean fish weight, which was calculated with a special software based on the daily feed amount).

The weight difference is the difference between the predicted weight when the sample was taken and the predicted weight when the starvation began, we believe this difference contains the fish growth information in the period between the last sample was taken and the fish were starved). The mean slaughtering weights of these 24 cages were regressed against these three variables (sample, sea temperature, weight difference). The estimates of the linear regression parameters are shown in figure 3.5. Sea temperature turned out to be insignificant and thus was omitted. Although the weight difference and sample are significant variables and have positive relationship (positive coefficient) with response, these two explanatory variables are not adequate to explain the variance of response ($R^2 = 0.55$). Figure 3.6 shows the plot of predicted slaughter weight against true slaughter weight.

	Estimate	Std. Error	Pr(> t)
intercept	2.0652	0.6719	0.00599 **
weight difference	0.5151	0.2234	0.03199 *
sample	0.5071	0.1620	0.00526 **

Figure 3.5: Estimates of parameters in the linear regression model.

With a good model with strong predictive ability, the blue dots should fall closely on the 1:1 solid black line. The figures shows a poor prediction ability of our linear regression model. A little tip is to standardize the sampling procedure, for example, taking the samples at a fixed time before slaughtering or if the slaughtering time is hard to predict, then drawing the samples just before starving with fixed sample size. If more data could be collected and provided (for example not only the mean weight, but also the individual weight of the samples), a more powerful statistical forecasting tool can be built, for reference see (LAS⁺11).




Figure 3.6: A plot of predicted slaughter weight against true slaughter weight (blue points), the solid black line is the 1:1 predicted weight and true weight.

3.3.5 Simulation study

We have developed two interactive web applications with **Shiny**, a **R** package. These web applications visualize the distribution of salmon weights.

Suppose we have a cage of 2000 salmon, and we know that the weights of them are normal distributed with mean being 5 kg and coefficient of variance (CV) equal to 0.1. Now we need to know how many fish in the cage fall in weight class 4-5 kg. The first web application was built for this purpose. Figure 3.7 shows a screenshot of the app for visualizing the distribution of number of salmons for a given weight class. On the right hand side of the figure we can see there is approximately 44% to 52% of fish fall in weight class 4-5kg, with maximum likelihood being 48%. We can easily check the density of other weight classes by changing the parameter values through sliding the bars on the left panel.

The **R** code that creates this web application can be found in appendix. To run the application, copy the code into **Rstudio** and run it.

We were shown that the salmon weight distribution in a cage is not always normal and typically is a bimodal distribution.

The two different modes may result from that the weights of female and male fishes in the cage have different probability density function. Therefore, we built another web app to visualize how the different weight distributions for males and females render a mixture distribution with two peaks. Figure 3.8 shows a screenshot of this web application.

On the left panel we set the mean of normal distribution for male weights to be 5 kg and for female it is 4 kg. We put the coefficient of variance for both the males and

Distribution of number of salmons for each weight class




Figure 3.7: A screenshot of a web application for visualizing the distribution of number of salmons for a given weight class.

females on 0.1 and the proportion of each gender in this cage is half-half. Then we can look at the density function for males (blue dashed line), females (red dashed line) and a mixture of both (black solid line) on the right hand side of the figure. Similar as the first app, we can change the parameter values on the left panel and get a flavor graphically of the mixture distribution. The R code for the two apps can be found in the appendix.

To run the applications, copy the codes into `Rstudio` and run them separately.

3.4 Summary

This brief report contains an analysis of statistical issues connected with the evolution of the distribution function for weights of salmon in a pen. Already from the fundamental deterministic model for the growth of average weights, one can see that parameter estimation needs to be performed as late as possible in the evolution of the distribution function. No doubt samples are more costly at this stage, but the information contained in the samples gives a better estimate of the distribution at slaughter time.

We constructed a statistical model to forecast, with uncertainty, slaughtering salmon weight distribution.

We have implemented the model in two web based applications, for which we provide the **R** code.

Mixed normal distribution function




Figure 3.8: A screenshot of a web application for visualizing the mixture distribution of salmon weights for males and females.

3.5 Appendix

In this appendix we give the R code for the routines **Shiny 1** and **Shiny 2** described in the report.

Shiny 1:

```
library(shiny)

# Define UI for application that draws a histogram
ui <- shinyUI(fluidPage(

  # Application title
  titlePanel("Distribution of number of salmons for each weight class"),

  # Sidebar with a slider input for number of bins
  sidebarLayout(
    sidebarPanel(
      sliderInput("mean",
                  "Mean of biomass distribution:",
                  min = 0,
                  max = 10,
                  value = 5, step=0.05),
      sliderInput("CV",
                  "CV of biomass distribution:",
                  min = 0,
                  max = 1,
                  value = 0.1, step=0.01),
      sliderInput("size",
                  "Total number of salmons",
                  min = 0,
                  max = 200000,
                  value = 2000, step=1),

      sliderInput("class", "Weight class(kg):",
                  min = 1, max = 10,
                  value = c(4,5))
    ),

    # Show a plot of the generated distribution
    mainPanel(
      plotOutput("distributionPlot")
    )
  )
))

distFun <- function(mean, CV, size, class){
  nSim=1000
```

```

population=vector()
classPercent=vector()
for (i in 1:nSim){
  population = rnorm(size,mean=mean,sd=CV*mean)

  subPopulation=population[which(population>class[1]& population<=class[2])]

  classPercent[i]=length(subPopulation)/size}
return(classPercent)
}

# Define server logic required to draw a histogram
server <- shinyServer(function(input, output){
  output$distributionPlot <- renderPlot({
    # generate bins based on input$bins from ui.R
    hist(distFun(input$mean,input$CV,input$size,input$class),prob = TRUE,
          xlab="Proportion of the weight calss",
          main = "Density of the weight class",
          col="corn flower blue")
    lines(density(distFun(input$mean,input$CV,input$size,input$class)),
          col="red",lwd=2)}
  )})

# Run the application
shinyApp(ui = ui, server = server)

```

Shiny 2:

```

library(shiny)

# Define UI for application that draws a histogram
ui <- shinyUI(fluidPage(

  # Application title
  titlePanel("Mixed normal distribution function"),

  # Sidebar with a slider input for number of bins
  sidebarLayout(
    sidebarPanel(
      sliderInput("meanM",
                  "mean of biomass distribution for males",
                  min = 0,
                  max = 10,
                  value = 5,step=0.05),
      sliderInput("meanF",
                  "mean of biomass distribution for females",
                  min = 0,
                  max = 10,
                  value = 5,step=0.05),
      sliderInput("CVM",
                  "CV of biomass distribution for males:",

```

```

        min = 0,
        max = 1,
        value = 0.1, step=0.01),
  sliderInput("CVF",
    "CV of biomass distribution for females:",
    min = 0,
    max = 1,
    value = 0.1, step=0.01),
  sliderInput("p",
    "proportion of males:",
    min = 0,
    max = 1,
    value = 0.1, step=0.01)
),

# Show a plot of the generated distribution
mainPanel(
  plotOutput("bimodalPlot")
)
))

wMale <- function(z, meanM, CVM) {
  exp(-.5*(z-meanM)^2/(CVM*meanM)^2)
}

wFemale <- function(z, meanF, CVF) {
  exp(-.5*(z-meanF)^2/(CVF*meanF)^2)
}

wAll <- function(z, meanM, CVM, meanF, CVF, p) {
  p*exp(-.5*(z-meanM)^2/(CVM*meanM)^2)+(1-p)*exp(-.5*(z-meanF)^2/(CVF*meanF)^2)
}

geom.mean.w <- Vectorize(function(z, wfn=w, mu.theta=0, sigma.theta, ...) {
  integrand <- function(theta) {
    log(wfn(z, theta, ...))*dnorm(theta, mean=mu.theta, sd=sigma.theta)
  }
  exp(integrate(integrand, mu.theta-3*sigma.theta,
    mu.theta+3*sigma.theta)$value)
})

# Define server logic required to draw a histogram
server <- shinyServer(function(input, output) {
  output$bimodalPlot <- renderPlot({
    # generate bins based on input$bins from ui.R
    par(mfrow=c(2,1))
    from <- 0
    to <- 10
    curve(wMale(x, input$meanM, input$CVM), from, to, ylim=c(0,1), lty=2, col="blue",

```

```
      xlab="Biomass(kg)", main="Biomass density",ylab=NA,lwd=2)
      curve(wFemale(x,input$meanF,input$CVF),add=TRUE,lty=2,col="red",lwd=2)
      curve(wAll(x,input$meanM,input$CVM,input$meanF,input$CVF,input$p),
            add=TRUE,lwd=2)
      legend(8,1,c("Male","Female","Mixture"),lty=c(2,2,1), lwd=c(2.5,2.5),
            col=c("blue","red","black"),bty = "n")
    },height=800)
  })

# Run the application
shinyApp(ui = ui, server = server)
```




References

- (FAA⁺16) M. Fre, M. Alver, J.O. Alfredsen, G. Marafioti, G. Senneset, J. Birkevold, F.W. Willumsen, G. Lange, Å Espmark, and B.F. Terjesen. Modelling growth performance and feeding behaviour of atlantic salmon (*salmo salar* L.) in commercial-size aquaculture net pens: Model details and validation through full-scale experiments. *Aquaculture*, 464:268–278, 2016.
- (JSW85) Kristinsson J.B, R.L Saunders, and A.J. Wiggs. Growth dynamics during the development of bimodal length-frequency distribution in juvenile atlantic salmon. *Aquaculture*, 45:1–20, 1985.
- (LAS⁺11) Anders Løland, Magne Aldrin, Gunnhildur Högnadóttir Huseby Steinbakk, Bang Ragnar, and Jon Arne Grøttum. Prediction of biomass in norwegian fish farms. *Canadian Journal of Fisheries and Aquatic Sciences*, 68(8):1420–1434, 2011.
- (MR01) T. Mørkøre and K. Rørvik. Seasonal variations in growth, feed utilisation and product quality of farmed atlantic salmon (*salmo salar*) transferred to seawater as 0+smolts or 1+smolts. *Aquaculture*, 199:145–147, 2001.



Problem 4. Electrode paste softening and flow

PROBLEM PRESENTED BY: AASGEIR VALDERHAUG, ELKEM ASA

STUDY GROUP: JESSICA WILLIAMS, CAOIMHE ROONEY, COLIN PLEASE, MARK MCGUINNESS, MATHIAS SANDULESCU, MADS PETER SØRENSEN

4.1 The Challenge

ELKEM ASA asked the Study Group to model the process of segregation in the Söderberg electrode. Segregation refers to the separation of two originally intermingled phases of the electrode during softening, producing a pool of liquid material adjacent to solid material. This is highly undesirable, as it may lead in turn to loss of control of electrode slip and possible destruction of the furnace.

The Söderberg electrode is a continuously created carbon electrode used in a variety of smelting furnaces, where the power is provided by passing an electric current (250 kA) into the the furnace via the electrode. The current and the high temperatures in the furnace erode the electrode, which is continuously fed into the furnace at about 2cm per hour. A row of these electrodes can be seen in operation in figure (4.1).

At the top of the electrode, a solid paste is fed into a steel tube at room temperature. This paste consists of calcined anthracite (bits of coke, 4–20 mm in size) bound together by tar pitch (which also contains finer coal particles), with about 70% by volume of the mix being the coarse coke. The paste is a solid cylinder¹ about 1m in diameter at the top of the electrode, sitting in a 2m diameter metal tube, and it

¹The paste may also be fed in as blocks or briquettes; we were asked to focus on the cylinders case



Figure 4.1: Søderberg electrodes at work in a smelting furnace. Photo courtesy of the main webpage <https://www.elkem.com/carbon/about-elkem-carbon/>




Figure 4.2: A sketch illustrating a Søderberg electrode in a smelting furnace, with an enlarged view on the right. The coloured contours indicate temperatures, increasing as the electrode approaches the furnace below. The smaller arrows indicate the direction of continuous movement of the electrode.

softens over a wide temperature range near 60°C as it descends and approaches the hot furnace, as illustrated in the sketch in figure (4.2). When operating correctly, the softened combination of coke and pitch spreads laterally to fill the surrounding metal tube as it descends into the furnace. This is illustrated as the blue-coloured region in the right-hand sketch in figure (4.2). When the softened paste reaches temperatures near 500°C it bakes into a solid carbon mass that can conduct the electric currents driving the furnace.




Figure 4.3: Sketches of Söderberg electrodes, illustrating the difference between a properly operating electrode (left side), and an electrode afflicted by segregation (right side).

When the electrode is operating correctly, there is no segregation. That is, the pitch remains well-mixed with the more solid coke particles as the whole mixture or paste softens and expands to fill its container sideways. Our challenge is to model the segregation process. When segregation occurs, the symptoms include the existence of a shiny lake (low viscosity fluid) of melted pitch, surrounding a more solid core of material, as illustrated in figure (4.3). The paste does not soften properly. Temperatures at the surface of the lake are low compared to temperatures at the surface of a properly softening paste near where it meets the enclosing cylinder. Photos taken from directly above the electrodes give another view of a well-behaved electrode compared to a segregated electrode, in figure (4.4).

The consequences of segregation can be that the solid baked region at the bottom



Figure 4.4: Photos taken looking directly down on Söderberg electrodes, illustrating the difference between a properly operating electrode (left, number 6), and an electrode affected by segregation (right, number 2). The fins inside the steel tube provide extra grip for the moving electrode near the bottom. Number 2 has a square block fed in at the top rather than a cylinder. Number 6 has multiple slabs fed in at the top, which are partially slumped.

of the electrode breaks off and falls into the furnace, together with the pitch lake, possibly causing an explosion and damaging the furnace. We seek to understand what causes the separation of the pitch and coke phases, by modelling it. The first step in our model is conceptual, to consider the paste to consist of two phases, pitch and coke.

4.2 Previous Work

Söderberg electrodes have been studied in the past, and there is a continuing interaction between the Industrially Focused Mathematical Modelling group at Oxford University and Elkem.

A 1993 paper by Fitt and Aitchison ([FA93](#)) addresses ways to measure the effective viscosity of the electrode paste at different temperatures. This paper arose out of a Study Group with Industry problem brought to Heriot-Watt University in 1988 by Elkem a/s, Norway. The shape of the surface of the bulging or slumping paste is related to its viscosity. Slow flow equations are developed from the Navier-Stokes equations, and solved analytically in the small time approximation. Exact solutions are also developed near the walls. Numerical solutions confirm the usefulness of the predicted effective viscosity of the electrode paste. Segregation is not specifically addressed in this paper.

We are also aware of a 2017 report by Alissa Kamilova ([Kam17](#)), supervised in part by Rolf Birkeland of Elkem. Segregation is pinpointed as a matter of concern in this report, but the focus of the project is to develop a mathematical model that produces the correct melting profile for a healthy electrode. Velocities of coarse particles and the binder pitch material are assumed to be the same, so that a separation of one phase from the other is not allowed to occur. Segregation is then

associated with a failure of the paste to soften, rather than with relative movement separating the pitch phase from the coarse particle phase.

Kamilova begins with separate mass and momentum conservation equations for the coke and the binder phases. Darcy's law is used, and an effective pressure is introduced, similar to the usual mathematical treatment of a saturated soil. An energy equation is added since viscosity is strongly temperature dependent. Effective pressure is assumed constant and the model is reduced to an equivalent single-phase medium. Steady-state equations are solved numerically to find the shape of the slumping surface of the softening paste. A key result of the modelling is that the temperature of the air surrounding the upper part on an electrode is very important for controlling the softening process.

4.3 Early Developments

4.3.1 Heat Transport

We expect that a key factor in segregation is the heating of the paste, so we computed the critical radius of the electrode for the time to heat it through radially by conduction to match the timescale of heat flow due to the vertical velocity imposed on the electrode (advection).

So we match

$$u_s T_z = \frac{k}{\rho c_p r^2} (r^2 T_r)_r$$

where for the paste, $u_s = 2\text{cm/hr}$ is the imposed vertical velocity, $\rho \approx 1000$ is density, $c_p \approx 1000$ is the specific heat capacity, and $k \approx 3$ is the thermal conductivity. In this report, SI units are used unless indicated otherwise. A dimensional approach then gives

$$\frac{u_s}{L} = \frac{k}{\rho c_p R_1^2}$$

where R_1 is the radius of the paste where it is being introduced at the top of the electrode, and L is the height of the electrode paste over which significant temperature changes occur.

Rearranging and substituting values in then gives

$$R_1^2 \approx \frac{L}{2}.$$

So for $L = 0.5$, R_1 is also 0.5, and this is a typical value used for R_1 in practice.

If L is larger than 0.5m, this simple heat analysis suggests that the electrode is moving too fast for heat to conduct all the way in from the sides, and R_1 should be reduced if the electrode is to soften successfully.




Figure 4.5: Sketch of three electrode regions, two solid sections sandwiching a narrow central region containing low viscosity fluid.

4.3.2 Lubrication Flow

If we consider instead the interaction of pressure and viscosity, conservation of momentum considerations lead to Reynold's lubrication equation for the flow of a viscous fluid through a narrow gap. The sketch in figure (4.5) illustrates the concept behind this analysis — two solid electrode regions sandwich a relatively low viscosity region that is narrow. The upper region is cold and solid, the lowest region is hot and baked solid, and segregation may occur in the central region of vertical extent h say.

In the steady state, Reynold's lubrication equation is

$$\nabla \left(\frac{\rho h^3}{12\mu} \nabla p \right) = \rho(u_s - u_f)$$

and we scale pressure as $p \sim \rho g H \tilde{p}$ where H is the height of solid paste overlying the low viscosity region, and radius as $r \sim R_1 \tilde{r}$ where R_1 is the radius of the paste where it is fed into the top of the electrode.

These scalings, neglecting the velocity u_f , give a balance in this steady-state provided that

$$h^3 \sim \frac{12\mu u_s R_1^2}{\rho g H}.$$

This formula gives an estimate of how the thickness of the low-viscosity region depends on the radius R_1 . It is based on balancing the overburden pressure with the speed of flow of a viscous fluid through the narrow region of width h . We can use this to inform segregation, by positing that if h is too small, the solid phase of the paste with its larger pieces of coke cannot flow through the low viscosity region, and hence segregation occurs. Then the critical value of R_1 , below which segregation is

anticipated, is

$$R_1^{\text{crit}} = \sqrt{\frac{G^3 \rho g H}{12 \mu u_s}}$$

where G is a measure of grain size of the coke phase. Taking this to be about 1mm gives

$$R_1^{\text{crit}} \approx 0.3m .$$

This is comparable to, but a bit smaller than, currently used values of R_1 .

Note that this analysis is somewhat at odds with the thermal analysis above — larger values of R_1 here will give larger h values and should be less susceptible to segregation from a lubrication flow perspective; however, larger R_1 values will be less well-heated (less softened) in the time taken to pass into the furnace, which may increase the risk of segregation since softening and slumping is a characteristic of a healthy electrode.

This suggests a coupled model that includes both thermal and viscosity effects is required. We begin to address such a model in the next section.

4.4 A Fully Coupled Model

Further work done at Oxford during a short intensive workshop on Söderburg electrodes at the end of May in 2018 extended the work of Kamilova, and was informally written up as a report ([sVBT+18](#)) at the time of our Study Group. One of the authors, Colin Please, was able to provide us with a copy of this report, and we made use of the model developed there, in our Study Group. That model was still in the early development stage. In particular, we adopted the conceptual model, and we focussed on finding the thickness of the narrow mobile region.

4.4.1 Conceptual Model

The idea underlying the conceptual model in the report ([sVBT+18](#)) is that estimates of the Peclet number for heat flow in the electrode are large. Hence advective heat flow due to the downwards movement of the electrode is much larger than conductive heat flow over the entire electrode column. The new twist on this result is to consider the implication, that there may then be a relatively narrow vertical region or thermal boundary layer in which advective and conductive heat transport mechanisms are comparable. This region is expected to lie between two relatively solid regions, an upper region where the paste is cool and hence solid, and a lower region where the paste is hot enough to have baked solid. This is then the region within which significant temperature changes occur, and within which segregation may occur. It is also near the region where segregation is observed to occur.

4.4.2 Equations

We adopt an axisymmetric, cylindrical polar coordinate system with z vertical along the long axis of the cylinder and r the radial coordinate. We consider that heat travels via both diffusion and advection. The steady heat equation can then be written in terms of temperature $T(r, z)$ as

$$\rho c_p (uT_r + vT_z) = -kT_{zz}, \quad (4.1)$$

where ρ is the density of the paste, c_p is the specific heat capacity, and k is the thermal conductivity. The variables $u(r, z)$ and $v(r, z)$ are the paste velocity components in the radial and axial directions, respectively. Since the flow region is thin, we begin with the thin film equations expressing conservation of mass and momentum for Newtonian fluid flow in a lubrication squeeze film,

$$(\mu u_z)_z = p_r, \quad \frac{1}{r}(ru)_r + v_z = 0, \quad p_z = 0, \quad (4.2)$$

where $p(r, z)$ is the pressure within the paste, and μ is the viscosity, which is a function of temperature, T . We approximate the relationship between viscosity and temperature as

$$\mu = \mu_0 \exp\left(\frac{T - T_{\min}}{T_0}\right)^2, \quad (4.3)$$

where $T_{\min}$ is the temperature where viscosity is the lowest, and T_0 is the temperature difference between the cold cylinder blocks and $T_{\min}$ ². This form gives the desired variation of viscosity from high values at low temperatures, through low values at temperatures near $T_{\min}$, to high values at high temperatures where baking has occurred, as sketched in figure (4.6).




Figure 4.6: Sketch of electrode paste viscosities (Pa.s) versus temperature (°C)

²We could equivalently take T_0 to be the difference between the hot part of the furnace and $T_{\min}$

4.4.3 Boundary conditions

We define $z = 0$ to be the top of the baked region of the electrode, which is at temperature $T = T_{\text{hot}}$. At the top of the electrode, $z = H$, we take $T = T_{\text{cold}}$, the temperature of the cold cylinders of paste when they are added to the electrode.

We assume that once the paste has baked, it will adhere to the sides of the enclosing steel cylinder, which move with downward velocity u_s . The paste cylinders are fed in at the top at speed u_f from the top of the furnace, where u_f will be determined so as to keep the height of the cylinder tower at constant height H .

The cold cylinders are of radius R_1 and the baked electrode is of radius R_2 . The variable radius of the paste as it deforms from cold cylinders to baked electrode is denoted $R(z)$.

4.4.4 Non-dimensionalisation

We non-dimensionalise using

$$r = R_1 \tilde{r}, \quad z = \delta R_1 \tilde{z}, \quad u = \frac{u_s}{\delta} \tilde{u}, \quad (4.4)$$

$$v = u_s \tilde{z}, \quad p = \rho g H \tilde{P}, \quad T = T_{\min} + T_0 \tilde{T}, \quad (4.5)$$

where $\delta \ll 1$ is a dimensionless parameter. We obtain from eqn (4.1) a dimensionless temperature equation

$$\delta \text{Pe} (\tilde{u} \tilde{T}_{\tilde{r}} + \tilde{v} \tilde{T}_{\tilde{z}}) = -\tilde{T}_{\tilde{z}\tilde{z}}, \quad (4.6)$$

where $\text{Pe} = u_s R_1 \rho c_p / k$ is a Peclet number, relating heat transport by vertical advection to heat transport by radial diffusion.

We also obtain the dimensionless thin film equations

$$(\exp(\tilde{T}^2) \tilde{u}_{\tilde{z}})_{\tilde{z}} = \alpha \tilde{p}_{\tilde{r}}, \quad \frac{1}{\tilde{r}} (\tilde{r} \tilde{u})_{\tilde{r}} + \tilde{v}_{\tilde{z}} = 0, \quad p_{\tilde{z}} = 0, \quad (4.7)$$

where $\alpha = \rho g H R_1 \delta^3 / (\mu_0 u_s)$. Calculating the size of Pe , we determine that it is reasonably large, so we choose $\delta = 1/\text{Pe}$. Then the temperature equation becomes

$$\tilde{u} \tilde{T}_{\tilde{r}} + \tilde{v} \tilde{T}_{\tilde{z}} = -\tilde{T}_{\tilde{z}\tilde{z}}. \quad (4.8)$$

4.4.5 Three regions

We now consider a thin fluid region of height $\alpha^{-1/3}$, centred around $T = T_{\min}$, that is, $\tilde{T} = 0$, where the paste has low viscosity. Thus, our domain is separated into three regions:

1. a top, essentially solid region which moves down with the velocity of the cold cylinders,
2. a middle, boundary-layer region, which has low viscosity, and
3. a bottom, baked region, which has adhered to the sides of the steel cylinder containing the electrode, thanks in part to the fins on the inside of the steel cylinder, and moves down with the imposed velocity u_s .

Regions 1 and 3: cold cylinders and baked electrode

The fluid equations in regions 1 and 3 to $\mathcal{O}(\alpha^{-1/3})$ are given by

$$\tilde{u}_{\tilde{z}\tilde{z}} = 0, \quad \frac{1}{\tilde{r}} (\tilde{r}\tilde{u})_{\tilde{r}} + \tilde{v}_{\tilde{z}} = 0, \quad p_{\tilde{z}} = 0, \quad (4.9)$$

and, taking into account the no slip conditions at the top and bottom of the baked region, we have that $\tilde{u} = 0$, and that $\tilde{v} = -u_f/u_s$ at the top and $\tilde{v} = -1$ at the bottom. The temperature equation is thus

$$\tilde{T}_{\tilde{z}\tilde{z}} = -\tilde{v}_0 \tilde{T}_{\tilde{z}}, \quad (4.10)$$

where $\tilde{v}_0 = -u_f/u_s$ in region 1, and $v = -1$ in region 3.

Region 2: boundary layer

Within the boundary layer, we consider new variables s and η defined to be

$$\tilde{z} = A(\tilde{r}) + \alpha^{-1/3} \eta, \quad s = \tilde{R}(\tilde{z}) \tilde{r}, \quad (4.11)$$

where $R(z) = R_1 \tilde{R}(\tilde{z})$. Note that $A(\tilde{r})$ is the locus of the vertical centre of the low viscosity region, and we are expanding about this locus. We further rescale velocity as

$$\tilde{u} = \alpha^{1/3} \hat{u}. \quad (4.12)$$

Our $\mathcal{O}(\alpha^{-1/3})$ equations are

$$\tilde{T}_{\eta\eta} + \hat{u} A' \tilde{T}_{\eta} = 0, \quad \left(e^{\tilde{r}^2} \hat{u}_{\eta} \right)_{\eta} = \frac{\tilde{p}_s}{\tilde{R}(\eta)}, \quad \frac{1}{\tilde{R}(\eta)} (\hat{u}_s)_s + \tilde{v}_{\eta} = 0, \quad \tilde{p}_{\eta} = 0. \quad (4.13)$$

We find that the pressure within the boundary layer (a function of $\tilde{r}$ only) is given by

$$\tilde{p}(r) = 1 - \frac{A(\tilde{r})}{H}. \quad (4.14)$$

This is under the assumption that the pressure at the top of the baked region ($\tilde{z} = 0$ and $\eta = -1$) is constant and equal to $p = \rho g H$, and thus $\tilde{p} = 1$. We also assume that the pressure within the fluid decreases linearly with increasing η (hydrostatic approximation). Furthermore, the total pressure through any cross-section within the boundary layer is equal to the total pressure at $z = 0$, which is $\rho g H$. This can be written in terms of our boundary-layer variables as

$$\int_0^1 \left(1 - \frac{A(s)}{H}\right) (R(\tilde{\eta}))^2 s ds = \frac{1}{R_1^2}. \quad (4.15)$$



References

- (FA93) A.D. Fitt and J.M. Aitchison. Determining the effective viscosity of a carbon paste used for continuous electrode smelting. *Fluid Dynamics Research*, 11:37–59, 1993.
- (Kam17) Alissa Kamilova. Modelling the paste flow and segregation in high-performance søderberg electrodes. Technical Report CDT Mini-Project 1, InFoMM, University of Oxford, 2017.
- (sVBT⁺18) Åsgeir Valderhaug, Rolf Birkeland, Halvard Tveit, Jon Chapman, Ian Hewitt, Colin Please, Alissa Kamilova, Graham Benham, and Ferran Brosa Planella. A model of a søderberg electrode. A report summarising the model developed during a short workshop held in Oxford, 30–31 May 2018.



List of Participants

Ana Avdzhieva

Sofia University "St. Kliment Ohridski", Bulgaria

Hanan Batarfi

King Abdulaziz University, Jeddah, Saudi Arabia

Heinrich Brüggemann

Helse Stavanger HF, Norway

Jacqueline Christmas

University of Exeter, UK

Meret Gisvold Sandberg

SalMar, Kverva, Norway

Lois Marie Granskog

Trondheim, Norway

Markus Grasmair

NTNU, Trondheim, Norway

Arne Guttvik

SalMar, Kverva, Norway

Harald Hanche-Olsen

NTNU, Trondheim, Norway

Poul Hjorth

Technical University of Denmark, Kgs. Lyngby

Dietmar Hömberg

NTNU, Trondheim, Norway/WIAS, Berlin, Germany

Dmitry Ivanov

EFD Induction AS, Skien, Norway

Bianca Lichtblau

TU Berlin, Germany

Mark Joseph McGuinness

Victoria University of Wellington, New Zealand

Thomas Petzold

WIAS, Berlin, Germany

Torbjørn Ringholm

NTNU, Trondheim, Norway

Caoimhe Rooney

University of Oxford, UK

Ian Roper

University of Oxford, UK

Magnus Vikan Røsaeg

Salmar, Kverva, Norway

Thomas Roy

University of Oxford, UK

Mathias Sandulescu

NTNU, Ålesund, Norway

Mads Peter Sørensen

DTU Compute, Kgs. Lyngby, Denmark

John Tyssedal

NTNU, Trondheim, Norway

Aasgeir Valderhaug

ELKEM, Kristiansand, Norway

Siebe Bruno van Albada

NTNU, Ålesund, Norway

Christian Veje

University of Southern Denmark, Odense

Jessica Williams

University of Oxford, UK

Jiafeng Xu

NTNU, Trondheim, Norway

Cao Yihan

NTNU, Trondheim, Norway