Regular article

Probabilistic design of a molybdenum-base alloy using a neural network

B.D. Conduit [a], N.G. Jones [b], H.J. Stone [b], G.J. Conduit [c],*

[a] Rolls-Royce plc, P.O. Box 31, Derby DE24 8BJ, United Kingdom
[b] Department of Materials Science & Metallurgy, University of Cambridge, 27 Charles Babbage Road, Cambridge CB3 0FS, United Kingdom
[c] Cavendish Laboratory, University of Cambridge, J.J. Thomson Avenue, Cambridge CB3 0HE, United Kingdom

ARTICLE INFO

ABSTRACT

An artificial intelligence tool is exploited to discover and characterize a new molybdenum-base alloy that is the most likely to simultaneously satisfy targets of cost, phase stability, precipitate content, yield stress, and hardness. Experimental testing demonstrates that the proposed alloy fulfills the computational predictions, and furthermore the physical properties exceed those of other commercially available Mo-base alloys for forging-die applications.

© 2017 Acta Materialia Inc. Published by Elsevier Ltd. All rights reserved.

The contemporary approach to develop new materials is experiment driven trial and improvement [1]. This approach may take up to twenty years to design and verify a new material. The long lead time rules out designing new materials alongside products, instead forcing engineers to compromise products around the shortcomings of pre-existing materials. The opportunity to discover materials computationally has the potential to empower engineers to design optimal materials at the same time as new products [2], bringing materials into the heart of the design process. Previous approaches to design new materials on a computer include ranking compositions with a Pareto set [3–5], characterizing materials with a principal component analysis [6], robust design [7], and the orthogonal optimization of different properties [2,8-11]. However, these methods cannot simultaneously optimize the compromise between material properties and capture the deep correlations between composition and final properties. Therefore, in this paper, a new artificial intelligence tool [12] that can capture the full composition-property relationship is used to propose the new Mo-base alloy for forging die applications that is most likely to satisfy all target properties simultaneously.

Mo-base alloys offer exceptional strength at high temperature. This makes them suitable for refractory applications including fission and fusion reactors, rocket engine nozzles, furnace structural components, and forging dies. However, the next generation of forging applications will demand yet higher operating temperature requiring a new generation of Mo-base alloys. Existing Mo-base alloys such as MHC (1.1 wt% Hf, 0.1 wt% C, balance Mo), TZC (1.2 wt% Ti, 0.1 wt% C, 0.3 wt% Zr, balance Mo), TZM (0.5 wt% Ti, 0.02 wt% C, 0.08 wt% Zr, balance Mo), and ZHM (1.2 wt% Hf, 0.1 wt% C, 0.4 wt% Zr, balance Mo) [13] contain minimal strengthening precipitates, so there is an opportunity to optimize the content of HfC and other carbides in Mo-base alloys to improve strength at high-temperature. Critically, the effective exploitation of strengthening precipitates requires a firm understanding of the relationship that exists between the alloy composition and it phase stability, strength and cost; a multidimensional problem that is an ideal application of an artificial intelligence tool.

The first section of this paper outlines the artificial intelligence tool and specifies the chosen targets for the relevant material properties: cost, phase stability, HfC content, yield stress, and hardness. In the second section, the tool is used to propose the new Mo-base alloy that is most likely to exceed the design targets. The final section presents experimental results for the phase stability, HfC content, and hardness to verify the model predictions and demonstrate that the alloy has properties that surpass those of other commercially available Mo-base forging die alloys.

The goal of the neural network tool is to predict the composition and processing variables that are most likely to produce a material that fulfills the multi-criteria target specification. The tool and

* Corresponding author.

methodology follows the prescription developed in Ref. [12]. The tool first constructs a predictive model for each property as a function of the composition, which for the Mo-base alloys presented in this paper comprises of the elements {Nb, Ti, C, Zr, Hf, W, Mo}. The tool can then calculate the likelihood that a putative composition fulfills the target specification, so that it can search composition space for the alloy most likely to meet the target specification.

Materials must fulfill a wide ranging specification to ensure that they best meet the needs of their target application. The properties that were optimized in the design of the Mo-base alloys are shown in Table 1. With properties depending on contrasting physics, for each property a different source of data must be adopted, which are referenced in the tables. The cost per cycle – the effective cost per usage as a forging hammer, which must be minimized – is predicted using a model of the weighted commercial elemental prices. The alloys with the most suitable mechanical properties are expected to be those that possess a Mo solid solution containing only HfC and other carbide precipitates. The low diffusion constant in Mo alloys below 1500 °C [14,15] means that the phase stability and HfC content should reflect the likely room temperature condition of an as-cast alloy. The thermodynamic phase stability and HfC content is evaluated by a neural network trained on a database comprising of CAL-PHAD results, with the data sourced from the SSOL6 database [16,17]. The use of a neural network to predict phase stability dramatically speeds up the alloy optimization process as it is computationally less intensive than individual thermodynamic calculations. It is essential for forging-die alloys to be strong, particularly in compression, so both the yield stress and also the hardness must be maximized. However, the yield stress and hardness cannot be reliably calculated by computer modeling from first principles. Instead a database of experimental results for all of the properties as a function of composition is compiled from the sources referenced in Table 1 comprising of alloys in an as-cast condition and exclusively of the Mo solid solution phase behavior prescribed by the thermodynamic predictions. The scarcity of hardness data means that the neural network can be improved if it is supplemented with ultimate tensile strength data. The neural network formalism [12] can automatically identify the link between ultimate tensile strength and hardness (known to be approximately ×3 [18]) from common compositions, and then use the surplus ultimate tensile strength data at other compositions to guide the extrapolation of the hardness model.

After the database of material properties in Table 1 is compiled, a neural network model is trained on that data to predict the physical properties for a given composition. The form of neural network and approach to training follows that in Ref. [12] used to develop Ni-base superalloys. The design variables were the elemental concentration of {Nb, Ti, C, Zr, Hf, W, Mo}. Typically three hidden nodes gives the best fitting neural network. The neural network model predicted not only the expected value of the physical property but also the uncertainty associated with it, accounting for experimental uncertainty in the underlying data, the uncertainty in the extrapolation of the training data [19,20], and the uncertainty in the processing conditions of as-cast alloys.

In this approach, the individual material properties are converted into a single merit index that describes the likelihood that the material properties (**V**) satisfies the design criteria (**T**) is $L = \Phi[\Sigma^{-1}(\mathbf{V}-\mathbf{T})]$.

Table 1
The approach used to predict properties, the property targets, number of experimental points used to train neural network models, and references for the data are shown.

Property	Target	Approach	Data
Cost	< 52$/cycle	Physical	[31–35]
Phase stability	>81 wt%	CALPHAD	[16,17,36-40]
HfC content	>1 wt%	CALPHAD	[16,17,36-40]
Yield stress at 1000 °C	> 398MPa	Neural net	212 [41,42,42-51]
Hardness at 1000 °C	> 1908MPa	Neural net	740 [41,42,42-66]

Here ,Φ is the multivariate cumulative normal distribution function and Σ is the covariance matrix [21]. Combining the individual property likelihoods enables an estimate to be made of the likelihood that the alloy will fulfill the whole specification. Critically, this overall likelihood will be much lower than that of an individual property target being met. For example, for a five-part specification, if the material has a 50% likelihood of fulfilling each design criterion, the overall likelihood that it simultaneously fulfills five criteria is $0.5^5 \approx 0.03$, so 3%. It is therefore crucial that the likelihood of the material meeting the conformance specification is maximized. The use of likelihood also allows the tool to explore and select the ideal compromise between material properties, which is inaccessible with methods that do not account for likelihood, such as a principal component analysis [6] and robust design [7]. Similarly, the design tool may interpolate between experimental data, exploring more compositions than would be accessible by an experimentally driven search. Using a neural network to interpolate allows us to capture deeper correlations than linear regression methods such as those used in principal component analysis [6].

As well as predicting material properties, the tool must simultaneously optimize them against the set targets. Previous optimization techniques included running over a pre-determined grid of compositions, and then sieving them with trade-off diagrams [10], or a Pareto set [3–5]. However the expense of these methods scales exponentially with the number of design variables rendering them impractical. Another approach is to use genetic algorithms [22,23], but this approach is not mathematically guaranteed to find the optimal solution [24,25], and it displays poor performance in high dimensional problems [24,25]. Here, we maximize the logarithm of the likelihood $\log(L)$ to ensure that in the region where the material is predicted to not satisfy the specification the optimizer runs up a constant gradient slope that persistently favors the least optimized property. We explore the high-dimensional composition space with a random walk which uses a step length that is comparable to the accuracy with which a material could be manufactured, this is 0.1 wt% for the entire composition excluding the possibility of microsegregation. The tool typically search over ~10^8 sets of design variables in ~1 h to explore the space and search for an optimal material.

With the neural network tool in place it is now used to design a new Mo-base forging die alloy. Once, designed, the properties of the alloy are subsequently verified by experiment. The goal is to design a new Mo-base alloy that offers both improved high-temperature hardness and concomitant greater lifetime with lower in-service costs at ~ 1000 °C. This case study not only serves as an independent test of the alloy design approach, but moreover leads to an alloy with properties that exceed those of other, commercially available Mo-base alloys.

The first step to design an alloy is to set the target specification. This is shown in Table 1 and compared with commercially available Mo-base alloys in Fig. 1 (a). The alloy should be cheaper than the previous cheapest Mo forging alloy, TZC, at 52$cycle^{-1}. To avoid forming deleterious phases that could weaken the alloy it must have good phase stability, defined as the concentration of the Mo-base solid solution rather than other deleterious phases, comparable or better than previous Mo alloys of 81 wt %. At the same time, the Mo alloy should be strengthened by HfC and other carbides, so there should be at least 1 wt % HfC precipitate formation. The yield stress should be greater than 398 MPa at 1000 °C, that of the best alloy available, ZHM. The alloy should also have a hardness higher than the highest of the Mo alloys, MHC, of 1908 MPa at 1000 °C. These targets mean that the new Mo-base alloy will have properties superior to those of any commercially available alloy. Neural network models for the cost, phase stability, volume fraction of the reinforcing phase, yield stress, and hardness are trained using data from the references in Table 1. The neural networks will then be used to optimize the