Mechanics of Materials 124 (2018) 118–131

Contents lists available at ScienceDirect

Mechanics of Materials

journal homepage: www.elsevier.com/locate/mechmat

Research paper

Experimental and computational issues for automated extraction of plasticity parameters from spherical indentation

J.E. Campbell, R.P. Thompson, J. Dean, T.W. Clyne*

Department of Materials Science & Metallurgy, Cambridge University, 27 Charles Babbage Road, Cambridge CB3 0FS, UK

ARTICLE INFO

Keywords:
Indentation
Plasticity
Constitutive law
Anisotropy
Friction

ABSTRACT

Software packages are being developed for automated extraction of plasticity parameters from indentation data (primarily load-displacement plots, although residual indent dimension data are also likely to be useful). Their design must be closely integrated with the associated experimental measurements. The procedure involves iterative FE simulation of the penetration of a spherical indenter into a sample, with automated convergence on a best-fit set of parameter values characterizing the yielding and work hardening response of the material (in a constitutive law). This paper outlines the main issues involved in optimization of experimental conditions and model formulation. Illustrative experimental data are presented from extruded rods of 5 metallic materials. Experimental issues include the dimensional scales of the indenter radius, R, and the depth of penetration, δ, with δ/R (the "penetration ratio") being of particular significance. A brief study is presented of the potentially conflicting requirements of deforming a volume large enough to represent the response of the bulk and having a value of δ/R that creates plastic strains in a range that will adequately capture the work hardening response. A key conclusion of this study is that a "mid-range" indentation facility is likely to be optimal, with a load capability of at least a few kN, able to create δ/R values up to ~40%, with R ~0.5–2 mm. Other experimental issues include displacement measurement techniques, calibration of machine compliance and the possibility of material anisotropy (due to crystallographic texture). Issues related to formulation of the FE model include specification of the domain and mesh, selection of the constitutive plasticity law and simulation of interfacial friction. The convergence algorithm used is also described.

1. Introduction

There has been increasing focus over the past decade or two on obtaining (true) stress-strain curves (well beyond the elastic limit) from outcomes of instrumented indentation experiments (mainly load-displacement plots, although residual indent shapes can also be used). Since these stress-strain curves are regarded as prime indicators of the plasticity characteristics of a material, and indentation is a much more versatile and convenient procedure than conventional uniaxial testing, this quest has a strong motivation. The approaches used fall into two main categories. Many studies (Taljat et al., 1998; Herbert et al., 2001; Basu et al., 2006; Kang et al., 2006; Pelletier 2006; Guelorget et al., 2007; Xu and Chen 2010; Hamada et al., 2012; Hausild et al., 2012; Pathak and Kalidindi 2015) have sought to identify analytical formulations that can be applied to the experimental data. This has obvious attractions, since such a formulation, even if involving relatively complex expressions and algorithms, would allow rapid extraction of the stress-strain curves via a well-defined path. Unfortunately, the stress and strain fields beneath an indenter, even one with a simple shape such as a sphere, are complex and change with penetration depth, making it very difficult to identify realistic analytical relationships. The prospects for this approach, certainly in terms of having a robust procedure that can be applied to a wide range of materials, are not promising.

The alternative approach (Dao et al., 2001; Bolzon et al., 2004; Bouzakis and Michailidis 2004; Bouzakis and Michailidis 2006; Pelletier 2006; Guelorget et al., 2007; Heinrich et al., 2009; Dean et al., 2010; Bobzin et al., 2013; Patel and Kalidindi 2016; Dean and Clyne 2017) is to use FEM modeling to (accurately) capture these evolving stress and strain fields, with the challenge then being to establish the stress-strain curve most closely consistent with measured indentation outcomes. This is a major challenge, but the approach is conceptually transparent, rigorous and simple (which cannot be said of the first type of methodology). However, its wide implementation has been inhibited by the need to carry out FEM modeling runs that are specific to each individual case, and also by uncertainties about how to converge on the

* Corresponding author.
E-mail address: twc10@cam.ac.uk (T.W. Clyne).

https://doi.org/10.1016/j.mechmat.2018.06.004
Received 24 March 2018; Received in revised form 24 May 2018; Accepted 12 June 2018
Available online 19 June 2018

"best fit" stress-strain curve and how to assess the confidence that can be placed in it.

For a material with a given (uniaxial) stress-strain curve, assumed to be applicable to deviatoric (von Mises) components of stress and strain for multi-axial situations, FEM can readily be used to predict the load-displacement plot (and residual indent dimensions). This can be done for any given indenter shape, provided that the important boundary conditions (potentially including the effects of friction between indenter and sample) can be established. However, the inverse problem of inferring the stress-strain relationship from such a load-displacement plot is much more challenging, with considerable scope for ambiguity, such as different stress-strain relationships giving effectively the same load-displacement plot.

For both types of approach, it has been recognized (Futakawa et al., 2001; Bucaille et al., 2003; Capehart and Cheng 2003; Chollacoop et al., 2003; Cheng and Cheng 2004; Ma et al., 2012) that there may be advantages in obtaining more comprehensive sets of experimental data. For example, doing repeat runs with indenters having different shapes has often been proposed, and indeed it is logical that this should be helpful, since the way that the stress-strain curve influences the indentation outcomes will be different with different indenter shapes. It has occasionally been suggested that simply using different indenter sizes may also be helpful, but this is unlikely to create benefits, since the stress and strain fields beneath an indenter are scale-independent. For example, the fields created by penetration of a sphere to a depth corresponding to, say, 10% of its radius are identical for radii of, say, 10 μm and 10 mm. The absolute value of the load at this point will be 10^6 greater for the latter case, while the penetration will be 10^3 greater, but the information being provided about the stress-strain response of the material is the same, provided the volume being interrogated is in both cases large enough to be representative of the bulk response.

The main requirement now, in order for procedures (and dedicated software packages) to become widely accepted and employed, is clear identification of the factors that affect sensitivities and efficient convergence on "correct" solutions for inferred properties. There are several key issues, concerning both experimental procedures and computational formulation. Some of these, including the development of algorithms for convergence on best fit parameter combinations, have been addressed by Isselin et al. (2006), while Karthik et al. (2012), among others (Giannakopoulos and Suresh 1999; Taljat and Pharr 2004), explored the influence of friction, concluding that it has a significant effect at penetration ratios above about 20%, particularly on the residual indent shape. Other workers (Sun et al., 1999; Ullner et al., 2010; Van Vliet, Prchlik et al. 2011) have drawn attention to the significance of machine compliance in the context of indentation load-displacement data. The present paper is aimed at examining all of the main issues in some detail, including the relationships between the experimental procedures and the numerical simulations. This is done using a wide range of experimental indentation data, illustrating how they are used in an automated way within software packages to obtain the values of parameters in constitutive stress-strain laws.

2. Experimental issues

2.1. Choice of indenter shape

There are several powerful motivations for using spherical indenters. One of these is that, since it is not a self-similar shape, the stress and strain fields change qualitatively as penetration takes place. Hence, the information being obtained over different depth ranges is analogous to carrying out separate tests with different indenter shapes (reducing the likelihood of different stress-strain curves giving very similar load-displacement plots). This point has been clarified previously (Dean and Clyne 2017).

There are also more practical motivations. One is that a sphere is much less prone to becoming damaged than are shapes having edges or points, and is also easier to specify and manufacture. Spheres (of WC-based cermets, with hardness and stiffness values high enough for most purposes), having diameters in the preferred range of about 1–4 mm (see Section 2.2 below), are cheap and readily obtained. There is also reduced risk with spheres of encountering the computational problems that are often associated with simulation of behavior in regions of high local curvature (edges or points).

Finally, at least with (approximately) isotropic materials, a spherical indenter allows the FEM modeling to be radially symmetric (2-D), which is not possible with many shaped indenters. The potential need for very large numbers of iterative FEM runs makes this a more significant issue than it would be under most other circumstances. All of the work described in this paper relates to use of spherical indenters.

2.2. Length scale effects

It is important, when the objective is to extract bulk properties, to indent on a suitable scale, while retaining the key advantages of being able to test small, flat samples, to carry out point-to-point mapping of properties etc. In particular, the volume being interrogated must have a (stress-strain) response that is representative of the bulk. It is on this *meso*-scale (such that indents are large enough for representative material response, but small enough to allow small samples and mapping) that this type of work needs to be focused.

The minimum indent size for representative response depends on microstructure, but in many cases it will require deformation of an assembly of grains - at least about a dozen and preferably more. Only when such an assembly is being deformed is it possible to capture the influence, not only of the crystallographic texture of the material, but also of the way that cooperative deformation of neighboring grains takes place. This is likely to be affected, not only by texture, but also by factors such as the ease of grain boundary sliding. Simply taking the average of the load-displacement responses from indents made in a large number of individual grains will not even approximately capture the bulk response. (The same arguments would apply to carrying out conventional uniaxial tests on a set of single crystal samples having orientations representative of the texture of a polycrystal.) A crude rule of thumb might be that, viewed on the free surface, the indent should straddle at least "several" grains. Of course, the corresponding minimum indent diameter might range from below 1 μm to above 1 mm, but it will certainly be small enough in most cases to offer the attractions outlined above.

Grain sizes of around 100 μm or more are, of course, common. In general, therefore, indent diameters should be at least a few hundred μm. This does require relatively large indenters (~ mm dimensions) and therefore large loads (~ hundreds of N, or even several kN), which may be beyond the range of some indentation systems (but perhaps below the commonly-used ranges of some conventional mechanical testing systems). However, systems in this "intermediate" load range are in general easier and cheaper to construct and use than either of the other two types of system. Moreover, a relatively coarse scale of indentation minimizes the problems associated with surface roughness, oxide films, contamination etc.

There is also a further issue, which relates to the indenter penetration depth, δ, as a ratio to the indenter radius, R. It might be imagined that, while the load needed to penetrate to a given δ/R, and the stresses in the material, would depend strongly on the material (hardness), the strains would not. In fact, this is not really true, since materials with different work hardening characteristics tend to exhibit significantly different plastic strain fields (for a given δ/R). Furthermore, even if the peak strain is, say, 40%, the indentation response will be considerably more sensitive to much lower strain regions of the stress-strain curve, in which most of the plastic deformation takes place. This issue is examined quantitatively in Section 5.4.1.

2.3. Material anisotropy and homogeneity

Among factors of potential significance are anisotropy (different responses in different directions of testing) and inhomogeneity (different responses in different parts of the same sample). Of course, this also applies to other types of testing, including conventional uniaxial loading, but some differences are expected between that case and indentation. Firstly, indentation is expected to be more suitable for exploring point-to-point variations in properties - indeed, this is one of its key attractions. This can be taken down to a very fine scale, examining the responses of individual phases or thin surface coatings, although attention is being focused here on testing of volumes sufficiently large to ensure a representative "bulk" response. Nevertheless, the bulk response could vary significantly between different parts of a relatively large sample or component and indentation is expected to be suitable for exploring this.

Anisotropy is a slightly different issue. It can arise in several ways in different types of material, but the commonest cause, at least in metallic polycrystals, is crystallographic texture. Most metallic samples are textured, at least to some degree, and this will, in general, lead to both elastic and plastic anisotropy. Properties such as Young's modulus and yield stress often vary by up to 10%, and sometimes considerably more, in metallic samples. (Of course, the extreme case of strong texture is a single crystal and, depending on the crystal structure and a number of other factors, very large elastic and plastic anisotropies can be observed in these.) These variations may be significant compared with the level of precision being sought in the inferred plasticity parameters.

Conventional uniaxial testing will clearly provide information about the anisotropy of a material (provided samples can be obtained that are suitable for testing the material in different directions, which may be difficult in some cases). It is not, however, immediately clear whether the indentation response of an anisotropic material will be significantly different in different directions. Certainly both stress and strain fields are much more multi-axial than in conventional testing, particularly for relatively deep penetration. The expectation might therefore be for volume-averaged responses for all indentation directions. However, this is not obvious and the present work encompasses an attempt to explore the issue.

The main way in which this is being done here is to focus on testing of extruded rod. Such samples can be anisotropic and/or inhomogeneous (on a fairly coarse scale). Of course, the grain structure, including the associated texture, will depend on several factors. In fact, extruded rods often exhibit stronger (and simpler) textures than, say, rolled sheet or plate. For the rods used in the current work, axial and radial samples were obtained for uniaxial (compression) testing.

3. Experimental procedures

3.1. Sample preparation

The extruded rods were all obtained from commercial suppliers. They were aluminum ("Dural"), brass (α/β), copper (OFHC) and a low C steel, with respective diameters of 20 mm, 16 mm, 25 mm and 20 mm. A further material (with a low hardness) was produced by machining Cu rod down to 16 mm and annealing it (in a sealed ampoule), to induce recrystallization. For uniaxial testing, samples in the form of cubes of side 5 mm were produced by electrical discharge machining (EDM). Surfaces for indentation were in most cases polished to a finish of 1 μm, using an automatic system to ensure that the polished face was accurately parallel with the base. Particular care was needed with the annealed copper, which was very soft. Most surface preparation procedures tend to leave a work-hardened layer, although this is naturally more noticeable for soft materials. In the present case, indentation was relatively deep (~200–400 μm), so a layer a couple of μm or so in depth would have little effect. However, this would not necessarily be true for a layer 5–10 μm deep, which could readily be produced on a very soft material. For the annealed Cu, therefore, prolonged and careful polishing on a 0.25 μm cloth was employed.

The grain structures of the rods were all examined, using optical and scanning electron microscopy. The grain sizes varied appreciably, within and between the different rods, but they were all no larger than about 100 μm (in transverse section). For the indenter size used (see Section 3.3), the deformed volume thus contained a large number of grains in all cases.

3.2. Uniaxial testing

Conventional uniaxial testing (in compression) was carried out at room temperature (22 °C ± 2 °C), using an Instron 3367 screw-driven testing machine. A load cell with a capacity of 30 kN was employed. Testing was carried out under displacement control, at a rate of 10 μm s^{-1}. Since they were 5 mm thick, tests took about a minute and the strain rate was of the order of 10^{-3} s^{-1}. This is taken to be a quasi-static rate (such that a very similar material response would be obtained if it were either increased or decreased quite substantially), confirmed via a series of trials for all of the materials used in the current work. Several repeat tests were carried out in all cases. Both stress and strain levels were converted from nominal to true values for comparison with indentation-derived curves.

Samples were compressed between rigid (hardened steel) platens, using MoS_2 lubricant to minimize barreling. In practice, there will always be at least some friction between sample and platen, which could possibly have an effect on the stress-strain curve (cause a rise in the apparent work hardening rate), particularly at relatively high strains (> ~10%). However, in the current work this effect was probably small. The displacement was measured using an eddy current gauge having a resolution of about ± 0.25 μm. It was attached to the upper platen and actuated against the lower one. Any error arising from (elastic) deformation of the platens is unlikely to be significant and there was no need for any compliance calibration.

3.3. Indentation testing

The indentation testing was carried out using the same testing machine (and eddy current gauge) as for uniaxial testing, and also similar displacement rates. A spherical indenter of radius 1 mm was employed, made of a WC-Co cemented carbide (cermet). This sphere was located in a matching recess in a steel housing, where it was secured by brazing. The set-up is depicted in Fig. 1. Compliance calibration was needed, since it's important in work of this type that the displacement data should be obtained under conditions that correspond closely to those being simulated in the model. The compliance was measured by pushing the indenter into a matched recess (about 500 μm deep) in a 5 mm thick plate of alumina, the recess having been created via abrasive rotational honing with the same type of indenter as that used in the tests. After a short bedding-down regime, this gave a linear plot (i.e. constant compliance) and this gradient was subtracted from indentation load-displacement plots. (The contribution to the compliance from the alumina plate itself is considered to be negligible.)

This calibration coped with the compliance of the housing, including the braze layer between indenter and housing, and also that of the top half of the indenter. (It would also cope with the compliance of other parts of the loading train, in cases for which that was significant). However, it did not compensate for the (elastic) deformation of the bottom half of the indenter, which could be significant in the early stages of indentation (when the contact area is small and the stresses and strains in the indenter could be relatively large). This part of the indenter was therefore included in the modeled domain, as shown in Fig. 1. This has the minor advantage that stresses in the indenter are being monitored, so it is possible, for example, to check on whether there might be any danger of it undergoing plastic deformation.

The compliance measured in this way was fairly small

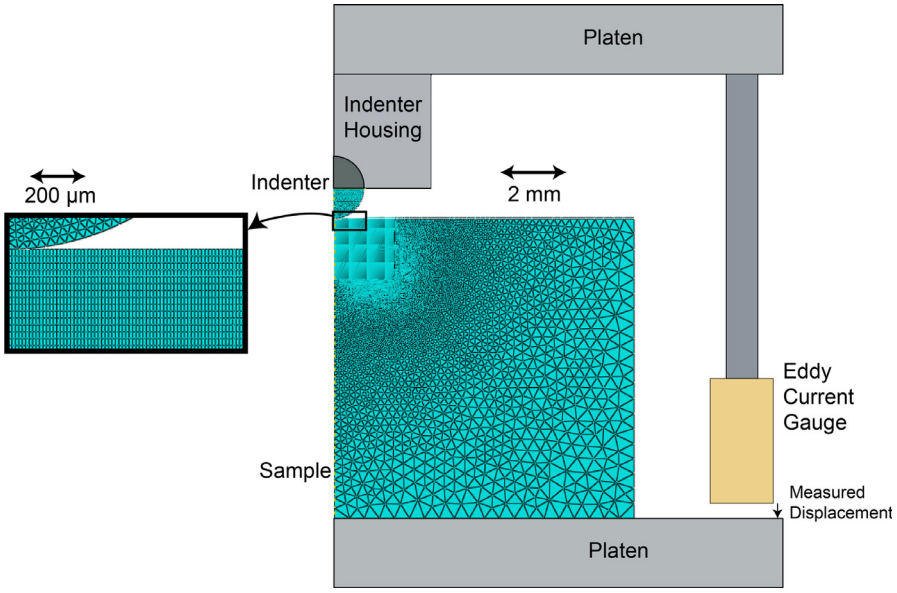

Fig. 1. Schematic of indentation testing set-up, showing the mesh used in the FE model.

(~5 μm kN^{-1}), so this was only a relatively minor correction. Other set-ups could be much more compliant, making this correction critical to the successful extraction of reliable property parameters.

3.4. Residual indent topography

A Taylor Hobson (Talysurf) profilometer (i.e. a contacting stylus), with a wide-range inductive gauge and 2 μm radius cone recess tip, was used to measure residual indent profiles. Scans were carried out in two perpendicular directions, both through the central axis of the indent. The height resolution of these scans is about 2 μm. Tilt correction functions were applied to the raw data, based on the far-field parts of the scan being parallel. The average profile from the two orthogonal scans was used for comparison with predicted profiles.

4. Computational issues

4.1. Constitutive plasticity laws

4.1.1. Background

For any approach involving iterative simulation of a deformation process, the stress-strain relationship (material plasticity response) must be characterized via a (small) set of parameter values. Of course, for a single simulation, it would be possible to use an arbitrary set of stress-strain data pairs, but when the objective is to infer the optimal relationship consistent with obtaining a particular outcome, this leaves too many degrees of freedom for tractable convergence, so a functional form (involving a relatively small number of parameters) is required. In fact, several expressions are in common use, some of them described in the following sections. It should be emphasized that these are all purely empirical relationships. There have been many efforts to rationalize stress-strain curves in terms of microstructural features (and their effect on dislocation mobility, which is at least the primary factor determining the ease of plastic deformation in metals). However, the concept of predicting stress-strain curves on the basis of identifiable features has not proved to be workable and the formulations in use are simply based on empirical fitting to experimental data.

A further point worthy of note concerns the possibility of asymmetry between (uniaxial) stress-strain curves obtained in tension and in compression. Any difference between the two is in principle indicative of a dependence of yielding (and subsequent progression of plastic straining) on the hydrostatic component of the stress state. In general, while differences are sometimes observed, they are normally due only to experimental difficulties - often associated with barreling (compression) or necking (tension). Genuine asymmetry is very rarely observed, at least for metals. This is consistent with the incompressibility of metals, the nature of metallic bonding and the main mechanisms of plastic deformation. The shear stress needed to cause dislocation glide (and also deformation twinning, which is significant in some cases) is effectively independent of the hydrostatic stress and so the uniaxial stress-strain relationship is expected to apply under any stress state, provided attention is focused on the deviatoric (von Mises) stresses and strains. This is implicit in virtually all FEM modeling of metal deformation.

4.1.2. Ludwik–Hollomon

Arguably the most common stress-strain relationship, and one that has been in use over an extended period (Hollomon 1945), is that usually designated the Ludwik–Hollomon equation:

$$\sigma = \sigma_Y + K\varepsilon_p^n \tag{1}$$

where σ is the (von Mises) applied stress, σ_Y is its value at yield, ε_p is the plastic (von Mises) strain, K is the work hardening coefficient and n is the work hardening exponent. A power law of this type is consistent with strain hardening being mainly caused by the increasing density of dislocations (impairing their mobility as more jogs, tangles etc. are formed), but with a hardening rate that falls off with continued straining (as the dislocation density approaches a saturation level). The fact that no power law expression actually allows complete saturation (a plateau stress) can lead to problems with this formulation, particularly at high strain levels. It can be seen that there are 3 parameter values in the general case.

4.1.3. Voce

The Voce relationship (Voce 1948) is commonly expressed in this form:

$$\sigma = \sigma_s - (\sigma_s - \sigma_Y)exp\left(\frac{-\varepsilon_p}{\varepsilon_0}\right) \tag{2}$$

where σ is the (von Mises) applied stress, σ_Y is its value at yield and ε_p is the plastic (von Mises) strain. The stress σ_s is a saturation level, while ε_0 is a characteristic strain for the exponential approach of the stress towards this level. This formulation thus simulates the competition between dislocation creation and annihilation reaching equilibrium (no

Table 1
Values of elastic constants used as input data.

Material	Young's modulus, E (GPa)	Poisson ratio, ν (-)
Al	68	0.33
Brass	107	0.33
Cu (both)	117	0.33
Low C steel	210	0.3
Cermet	650	0.21

further hardening) at sufficiently high strain levels (Mecking and Kocks 1981; Estrin and Mecking 1984; Sainath et al., 2015), although in practice the approach to this condition often does not conform very well to an exponential curve. In this case also, there are 3 unknown parameter values.

4.1.4. Other relationships

A number of other stress-strain relationships have been proposed, some aimed primarily at non-metallic systems such as polymers and rubbers. Some, such as that of Ramberg–Osgood, are oriented towards cyclic loading and a focus on the transition between elastic and plastic deformation (Skelton et al., 1997). There are others, such as the Ludwigson relationship, as described, for example, in Samuel and Rodriguez (Samuel and Rodriguez 2005), that are effectively combinations of the above two, but involve more parameters and hence are likely to slow down convergence considerably. While there could be a motivation in some circumstances for exploring a wider range of formulations than the above two, in general the stress-strain curves of most metallic materials can be captured reasonably well using at least one of these two equations.

4.1.5. Strain range

There is also the issue of the range of strain over which a representation is required, or is likely to be reliable. In practice, it is not usually either necessary or viable for levels above about 25% to be considered, although indentation testing is much better suited than uniaxial testing (tension or compression) to the reliable generation of higher strains than this and, if there is interest in the material response in this regime, it should probably be studied solely via indentation. In fact, the uniaxial (compressive) stress-strain curves presented here may not be entirely reliable beyond strains of about 10%, since the apparent work hardening rate could be raised slightly in this regime by the effect of friction.

Of course, in addition to plasticity characteristics, there may be interest in how damage and failure (crack propagation) might occur when strains start to get relatively large. However, despite various aspirations in this direction, there are severe limits to the information of this type that might be obtainable via indentation testing.

4.2. FEM Formulation, meshing and boundary conditions

4.2.1. Mesh specification

An axi-symmetric FEM model was employed, based on the mesh illustrated in Fig. 1. There are about 5000 volume elements, all second order quadrilateral and/or triangular. The mesh was refined in regions of the sample close to the indenter, as shown. Sensitivity analyses confirmed that the meshes employed were sufficiently fine to achieve convergence, numerical stability and mesh-independent results. The complete sample was included in the simulation, with its lower surface rigidly fixed in place. In modeling the complete sample, contributions to the displacement caused by its elastic deformation, as well as plastic deformation, are fully captured. In fact, the sample thickness is not important beyond a depth of about 5 indenter diameters, since the axial stress - and hence the elastic strain - below that becomes negligible for typical cases. The lateral extent, beyond about 2 or 3 indent diameters, is also unimportant.

4.2.2. Effect of interfacial friction

The effect of interfacial friction is routinely simulated via a coefficient of friction, μ, such that sliding between the two surfaces requires a shear stress, τ, given by

$$\tau = \mu\sigma_n \tag{3}$$

where σ_n is the normal stress at the interface. The value of μ is expected to depend on the surface roughness of indenter and sample, and cannot be predicted *a priori*. It may also be noted that new sample surface is created during indentation. Modeling experience showed that the predicted behavior can be fairly sensitive to the value, particularly as the penetration ratio starts to become relatively large ($> \sim 10\%$). The value of μ was therefore regarded as adjustable, so as to allow improvement of the fit between experimental and predicted load-displacement plots. (As mentioned in Section 4.3.1 below, it could have been included in the set of parameters to be iteratively optimized during convergence, but this was not done in the current work and it probably isn't necessary in most cases: in practice, it was found that best fit was usually achieved with a value of around 0.2, suggesting that this is typical of the effect of friction during experiments of this type.)

4.2.3. Input and output data

The input data included the elastic constants of the materials examined, plus those of the cermet. All were assumed to be isotropic. The values employed are listed in Table 1. The Young's moduli were measured directly, using an ultrasonic resonance system, although in all cases the values obtained were close to those in standard handbooks. Poisson ratio values were taken directly from handbooks. The simulation runs were carried out under displacement control, so the output was predicted loads at a series of (~ 50) specified displacement values over the range concerned. The residual indent shape, and the surrounding fields of residual stress and plastic strain, were also predicted in each case.

An investigation has also been made into how the plastic work is distributed in terms of prior strain. After each increment of strain, for each volume element, the stress, incremental strain and prior strain are recorded. The work done during that increment is evaluated (= stress × strain × volume) and that increment of work is associated with the strain concerned. Expressed mathematically, the increment of work done in the jth volume element during the kth increment of strain can be written

$$\Delta W_{j,k} = \sigma_{j,k}\Delta\varepsilon_{j,k}v_j \tag{4}$$

Clearly, the work done during the kth strain increment is given by

$$\Delta W_k = \sum_{j=1}^{j=M} \Delta W_{j,k} \tag{5}$$

where the summation is over the total number (M) of volume elements, and the total work done is

$$W_{tot} = \sum_{k=1}^{k=T} \Delta W_k \tag{6}$$

with this summation being over the total number (T) of strain increments. The total strain range is divided into a number of sub-ranges (bins) and the work done within each bin is then evaluated after a binning operation. This can be expressed as

$$W_{bin,p} = \sum_{k=1}^{k=T}\sum_{j=1}^{j=M} (\Delta W_{j,k} f_{j,k,p}) \tag{7}$$

where $f_{j,k,p}$ is a function ascribed a value of 1 or 0, depending on whether the strain associated with the increment of work $\Delta W_{j,k}$ does or does not fall within the range of the pth bin.

4.3. Convergence procedure

4.3.1. Quantification of the goodness-of-fit

The procedure used in the current work to quantify the goodness-of-fit between predicted and target indentation outcomes was least squares regression. This is popular for optimizing a set of model parameter values, by quantifying the fit between the modeled values of a scalar variable and corresponding expected values (Riley et al., 2006). These are the parameter values that minimize the sum of the squares of the residuals, which are the differences between expected and modeled values of the variable.

For the current work, the main outcome is a load-displacement plot, the variable is the load (as a function of displacement) and the parameters are those in the selected constitutive law for material plasticity. It would also be possible to include in this set other parameters that influence the indentation process, such as the coefficient of friction, although that has not been done for the results presented here. It may also be noted that other outcomes could be used, either alternatively or additionally. For plasticity, the other outcome that is likely to be used is the residual indent shape (indent depth as a function of radial location), although again that has not been done in the current work.

More generally, when other material properties are being sought, there are several options concerning the outcomes that might be obtained from experiment. For example, in the recent work of Burley et al. (2018), aimed at evaluating a strain rate sensitivity parameter from ballistic indentation experiments, both displacement-time and residual indent shape datasets were used as outcomes. Having more than one outcome can complicate the convergence procedure, although, in that particular case (with just a single material property parameter to evaluate), simple linear scans in parameter space allowed rapid identification of optimum values.

The sum of the squares of the residuals, S, can be expressed:

$$S = \sum_{i=1}^{N} (P_{i,\mathrm{M}} - P_{i,\mathrm{E}})^2 \tag{8}$$

where $P_{i,\mathrm{M}}$ is the ith value of the modeled displacement (predicted by FEM) and $P_{i,\mathrm{E}}$ is the corresponding experimental (target) value. The value employed for N was around 50. Perfect fit will lead to a value of zero for S. Since S is dimensional, it has units and its magnitude cannot be used to give a universal indication of the quality of the fit. For this purpose, the quantity S_{red}, a "reduced sum of squares" is used, defined by

$$S_{\mathrm{red}} = \frac{\sum_{i=1}^{N} (P_{i,\mathrm{M}} - P_{i,\mathrm{E}})^2}{N P_{\mathrm{av,E}}^2} \tag{9}$$

where $P_{\mathrm{av,E}}$ is the average of the experimentally-measured loads (across the range of displacements being used) and N is the total number of displacement increments. The parameter S_{red} is a positive dimensionless number, with a value that ranges upwards from 0 (corresponding to perfect fit). As a generalization, modeling that captures the material plasticity response reasonably well should lead to a solution (set of parameter values) for which S_{red} is less than, say, 10^{-3}. This effectively constitutes a health check on the solution - if, for example, no solution can be found giving a value smaller than, say, 1%, then this suggests that there can only be limited confidence in the inferred set of values. This could be due to experimental deficiencies and/or an inability to capture the behavior well with the constitutive law being used. In fact, during the work described here, solutions with S_{red} values around 10^{-4} or below were found in all cases.

4.3.2. Searching of parameter space

When there is no analytical relationship between experimental outcomes and parameter values, which is clearly the case here, the best-fit set of values must be estimated by iterative improvement, using a search algorithm. The algorithm used in the current work is the Nelder–Mead simplex search (Nelder and Mead 1965), which has been chosen in view of its robustness and adaptability, particularly with respect to noise. The Nelder–Mead search procedure used here is based on the implementation by Gao and Han (2012), and is built using the Scientific Python and Numeric Python packages (Oliphant 2007; van der Walt, Colbert et al. 2011).

For a model with m parameters, searching is within an m-dimensional parameter space, within which a *simplex* is defined. This is a *polytope* with $(m + 1)$ vertices (i.e. a triangle in 2-D, a tetrahedron in 3-D etc.). Each vertex corresponds to a particular combination of all of the m parameters in the set and the simplex covers a range of values for all of these. These points can be expressed as vectors (first rank tensors) in parameter space, designated $\boldsymbol{x}_1$, $\boldsymbol{x}_2$,…$\boldsymbol{x}_{m+1}$, each of which consists of a set of m parameter values. After each iteration (new set of FEM simulations), the objective is to "improve" the simplex by replacing the worst vertex (i.e. the one with the highest value of S) with a better point. The search for this better point is along a line in parameter space defined by the worst point and the centroid of the rest of the simplex, which is the average position of the remaining points (after removal of the worst point). The steps involved in the algorithm are described in the Appendix.

4.3.3. Convergence on the optimal solution

The algorithm is terminated once a specified convergence criterion has been met. This can be defined as a relative difference (commonly 10^{-4}) in S and/or $\boldsymbol{x}$ between successive iterations. In the present work, this requirement was applied to both S and $\boldsymbol{x}$. The number of iterations to achieve convergence depends on a number of factors, in addition to this criterion specification. These mostly relate to the way that the goodness of fit varies in parameter space, which in turn depends on several issues (including how well the stress-strain curve can be captured by the selected constitutive law). There may in some cases be a danger of converging on a local minimum. Difficulties can also arise from the presence of "plateau regions", where various parameter value combinations give very similar degrees of fit. In general, however, it has been found that the algorithm employed performs well in this context – see Section 5.3.

Overall, the computational operation is tractable in most cases. The procedure could, however, be facilitated if the starting values for the simplex were in an appropriate region of parameter space - i.e. if the initial trial values are fairly close to the best solution set. Otherwise, the number of iterations for convergence may go up significantly. One way of ensuring that these starting values are at least in the right part of parameter space would be to generate a pre-run matrix of simulations, and resultant load-displacement plots, and simply select the one giving the best fit (lowest S value). One drawback of this is that all such "pre-running" would need to be done for specified values of certain parameters that are not in the set being searched, notably the elastic constants of the material and the friction coefficient. It would also need to be done for at least two constitutive laws. However, provided the "pre-run" matrix is a relatively coarse one, it would not be a prohibitive operation to do this for several values of those parameters.

5. Illustrative results

5.1. Stress-Strain curves and effects of anisotropy

An initial check was made on the homogeneity of the rods, since any significant heterogeneity would clearly need to be taken into account. This was done using simple Vickers hardness measurements. The outcome is shown in Fig. 2(a), where it can be seen that, in all cases, there was little or no variation with radial location. Representative uniaxial

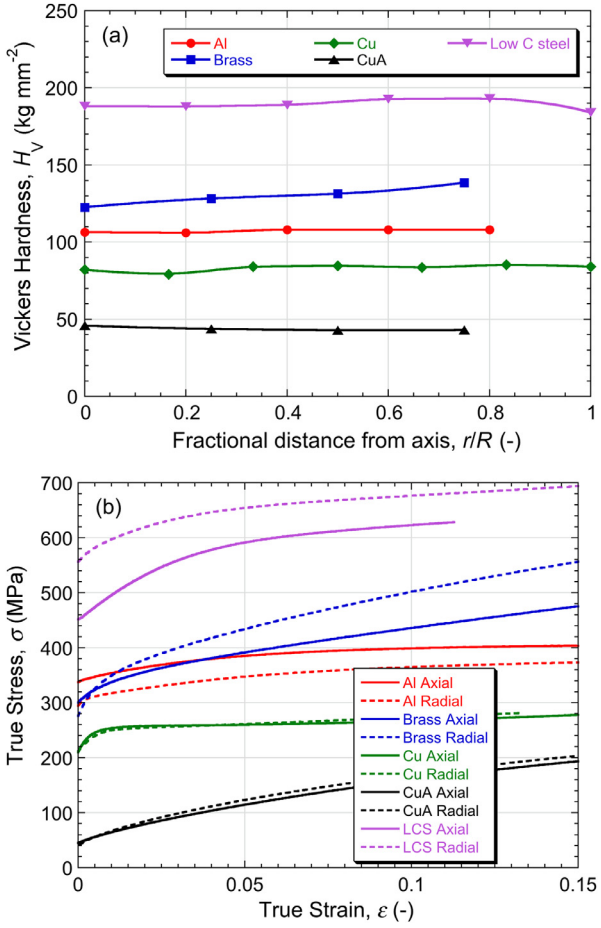

Fig. 2. (a) Vickers Hardness numbers, as a function of radial location on transverse sections and (b) stress-strain plots, from compressive loading in both axial and radial directions.

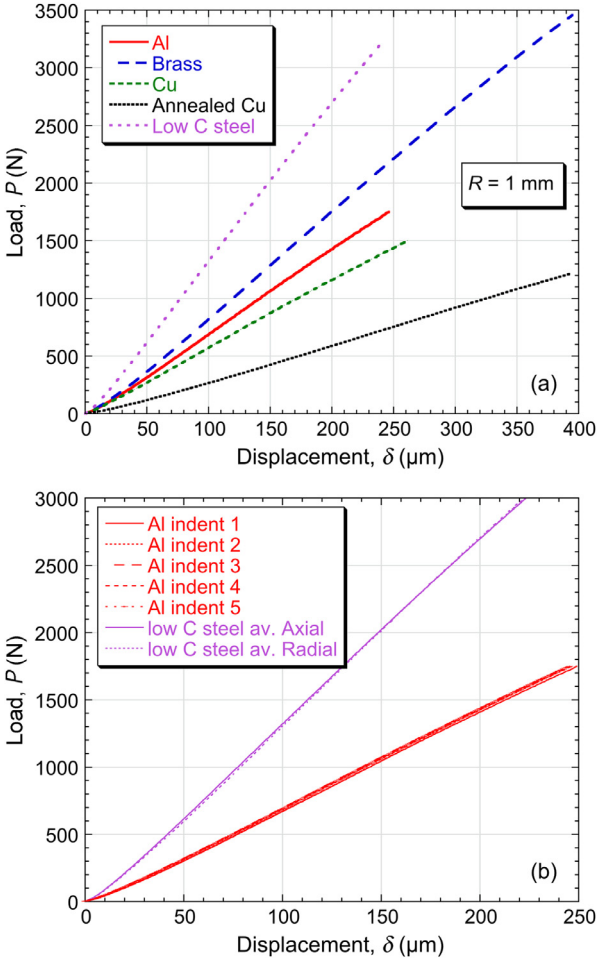

Fig. 3. Sets of indentation load-displacement data, comparing: (a) averaged plots from transverse sections and (b) individual plots for the Al (axial), together with axial and radial average plots for the low C steel.

stress-strain plots (in both axial and radial directions) are shown for all 5 metals in Fig. 2(b). For Cu (both as-received and annealed), little anisotropy is exhibited between these two directions. For the brass, the yield stresses are similar in the two directions, but the work hardening is significantly different. For the Al, the yield stresses differ slightly, but work hardening is similar. For the low C steel, both yield stress and work hardening differ. Furthermore, there is no systematic trend regarding which of the two directions is harder. It might be possible to rationalize these characteristics in terms of the nature and strength of the textures in each case, although this would require further characterization.

For present purposes, it should just be noted that this set of materials exhibits a wide range of plasticity responses, with yield stresses from ~50 to ~600 MPa and average work hardening rates (up to ~15% strain) varying from below 100 MPa (Al and as-received Cu) to well over 1000 MPa (brass). Also, for the annealed Cu, the flow stress rises by 400% over this strain range - i.e. work hardening is highly significant - whereas for the Al the relative increase is only about 20%. Moreover, the shapes of the work-hardening curves vary significantly. Such a set of curves provides an excellent test for the reliability of a methodology for their inference from indentation data.

5.2. Outcomes of indentation experiments

The main outcomes are load-displacement plots from indents made at various locations on transverse sections of the 5 rods. These were all more or less independent of radial position, which is consistent with the hardness data shown in Fig. 2(a). Representative plots are shown in Fig. 3(a). In addition to the excellent reproducibility of these data, it should be noted that, for any given material, the plots obtained by indenting radially were very similar to those obtained on transverse sections. Both of these points are illustrated by the plots shown in Fig. 3(b). In view of the significant anisotropy apparent in the uniaxial stress-strain curves (Fig. 2(b)) of some of these materials (particularly the low C steel, for which the indentation comparison is shown in Fig. 3(b)), this outcome confirms that indentation responses are highly-multi-axial. The implication is that indentation, perhaps particularly when carried out to high penetration ratios, will lead to inferred stress-strain curves that are averages of those obtained uniaxially in different directions. While this might be regarded as a limitation in some respects, it does mean that there need be no concern about the direction in which indentation is carried out.

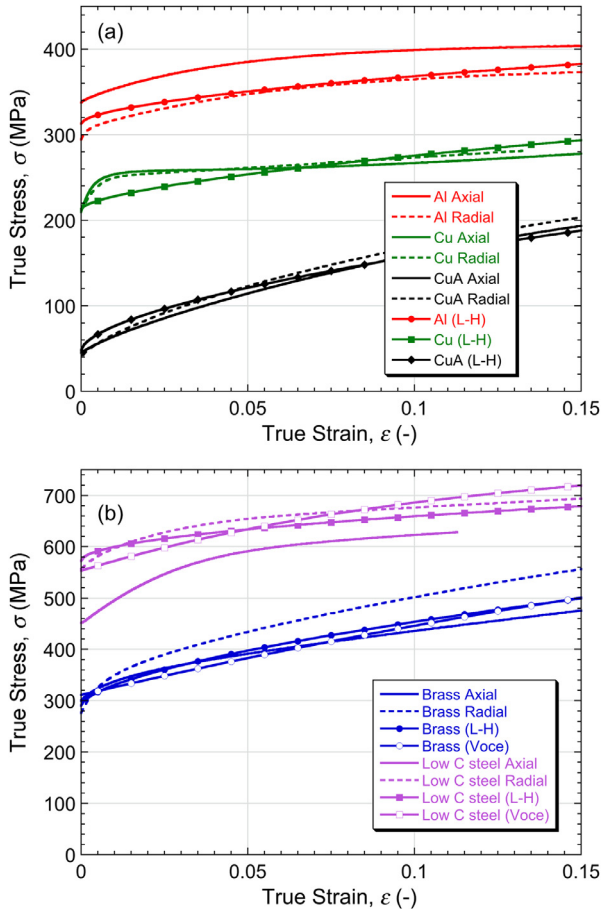

Fig. 4. Comparisons between experimental stress-strain curves (axial and radial directions) and corresponding plots of constitutive laws, using the parameter values obtained by iterative comparison between measured and modeled indentation load-displacement data - see Table 2. These comparisons are shown for (a) Al, Cu and annealed Cu, and for (b) brass and low C steel.

5.3. Inferred stress–strain curves

A comparison (for all 5 metals) is shown in Fig. 4 between uniaxial stress-strain plots (axial and radial) and those obtained via Nelder–-Mead convergence on best-fit parameter combinations. This has been done using both Ludwik–Hollomon and Voce expressions for all 5 metals. The plots in Fig. 4(a) show only L-H predictions, while those in Fig. 4(b) show both. It can be seen that the level of agreement is very good. It could be quantified in several ways, but in all cases the predicted curves are within about ± 5% of the average of the axial and radial uniaxial plots, over the complete strain range. This is certainly acceptable for most purposes and indeed it might be argued that uniaxial data can rarely be considered reliable to higher precision than this. In general, similar inferred stress-strain curves were obtained using the two constitutive laws, although, as can be seen for the low C steel (Fig. 4(b)), they do sometimes have slightly different shapes. It's certainly recommended that alternative expressions be explored, since there may be cases for which one or the other gives noticeably better consistency with experimental (indentation) data.

The convergence procedure generally required only a relatively small number (a few tens) of iterations, although of course this does depend on the level of accuracy being sought. Some features of the convergence are apparent in Fig. 5, where it can be seen how the movement through parameter space took place for one of these materials. The starting values were selected in a fairly arbitrary way and, as can be seen, they were not very close to the optimized values in this case. If they happened to be either very close to, or to differ substantially from, the optimized set, then that would affect the efficiency of the convergence. However, there would not, in general, be very much difference between the number of iteration required in the two cases. On the other hand, there will always be a requirement for more than just a few iterations, even if the starting positions are quite close (obtained, for example, by scanning a pre-run set).

All of the inferred parameter values are shown in Table 2. Also shown are σ_Y values from H_V numbers, obtained via the commonly-used approximation $\sigma_Y \sim H_V/3$ (with H_V, expressed in MPa, being averages of the values shown in Fig. 2(a)). As might have been anticipated, these are at least in the right general vicinity for the materials that exhibit relatively little work hardening (Al, Cu and low C steel). However, for the other two materials (brass and annealed Cu), which work-harden quite strongly, they are massive over-estimates. Of course, a yield stress obtained from a hardness measurement should never be regarded as better than semi-quantitative.

5.4. Sensitivity issues

5.4.1. Penetration ratio

A key issue concerns the penetration ratio (δ/R) necessary to sample the stress-strain characteristics of the material to sufficiently high levels of strain. Clearly, if the plastic strains created beneath the indenter do not extend beyond a few %, then the outcome will not be sensitive to the work hardening characteristics. Moreover, the levels of strain created at a given (δ/R) do depend on the work hardening characteristics. This is illustrated by the strain fields shown in Fig. 6, which are for 40% penetration into the as-received and annealed copper. It can be seen that, while the peak strains are ~150% for the as-received Cu, they are only ~60% for the annealed Cu. This is because the more pronounced work hardening of the latter leads to the strains becoming significant in a larger surrounding region, while limiting the peak levels close to the indenter. It can also be seen that the two residual indent shapes are very different, with considerably more “pile-up” in the material that does not work harden much. It follows that, for a material that exhibits pronounced work hardening, deeper penetration is advisable, in order to interrogate higher strain regions of the stress-strain curve.

There is, of course, more to this than just the issue of peak strain levels. The real requirement here is an insight into how much of the deformation (that influenced the indentation response) took place in different regimes of strain. Information of this type (obtained using the methodology described in Section 4.2.4) is presented in Fig. 7, with Fig. 7(a) and (b) showing fractions of the plastic work done, as a function of the strain at which it took place, in the two Cu materials, for 3 different (δ/R) ratios. Also shown (Fig. 7(c)) is a plot of the average strain at which the plastic work was done, for all of the metals, as a function of (δ/R). These average strains tend to be lower, at a given (δ/R), for metals that exhibit greater work hardening. A general conclusion of studies of this type is that, in order to be confident of ensuring an indentation outcome that is suitably sensitive to behavior over the strain rage of interest (up to ~20%), with no prior knowledge of the material plasticity characteristics, a penetration ratio of ~40% is recommended.

There is also interest in the stress fields generated during

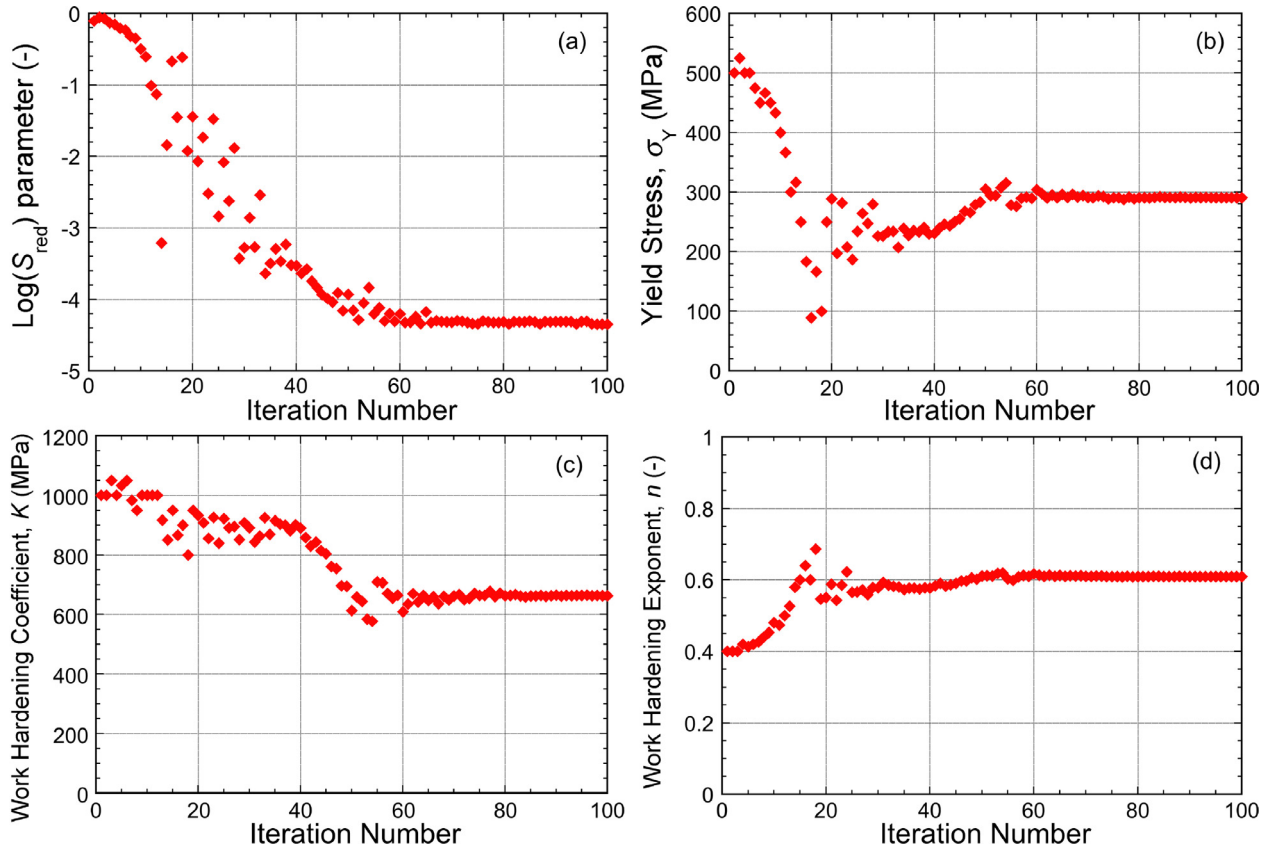

Fig. 5. Nelder–Mead convergence on an optimal (Ludwik–Hollomon) parameter set, targeting a load-displacement plot from indentation of a brass sample, showing the evolution with iteration number of: (a) the goodness-of-fit parameter, S_{red}, (b) yield stress, (c) work hardening coefficient and (d) work hardening exponent.

Table 2
Inferred plasticity parameter values.

Metal	H_V-derived yield stress, σ_Y (MPa)	Ludwik–Hollomon			Voce		
		Yield stress, σ_Y (MPa)	WH Coefficient, K (MPa)	WH Exponent n (-)	Yield stress, σ_Y (MPa)	Saturation stress, σ_s (MPa)	Characteristic strain, ε_0 (-)
Al	351	312.8	204.4	0.565	325.9	518.5	0.455
Brass	425	290.4	661.7	0.608	311.1	885.6	0.374
Cu	272	213.1	263.8	0.626	218.3	315.6	0.145
Cu (anneal)	143	47.9	428.6	0.590	55.9	358.4	0.232
Low C steel	618	571.8	274.9	0.497	553.1	778.1	0.112

indentation, partly in terms of the possibility of the indenter undergoing plastic deformation. An illustrative example is presented in Fig. 8, which shows the von Mises stress field, at a penetration ratio of 25%, for the low C steel. These stress levels correspond to the flow stress for the local level of plastic strain. In the immediate vicinity of the indenter, these levels are ~700–900 MPa, corresponding to plastic strains of several tens of % (see Fig. 4(b)). The stresses in the indenter are also worth noting. They reach levels approaching 2 GPa in this case - this is expected to be below the yield stress of the cermet, but is nevertheless a substantial stress. Use of a steel for the indenter, rather than a cermet, is not recommended, partly because such stress levels could cause plasticity and partly because the much lower stiffness (~200 GPa, compared with ~650 GPa for a cermet) means that displacements due to elastic deformation of the indenter will be larger and more likely to constitute a significant source of potential error.

5.4.2. Coefficient of friction

It is difficult, if not impossible, to directly measure a coefficient of friction under the conditions of an indentation test (during which new sample surface is being continually created). It's also difficult to vary the value in any systematic way - both indenter and sample are expected to be smooth and, since the contact pressure is high, it is not practicable to introduce lubricant. However, it is straightforward to vary the value employed in the model and explore the effect that this has. This has been done in the present work using both load-displacement data and residual indent profiles. The latter are particularly well-

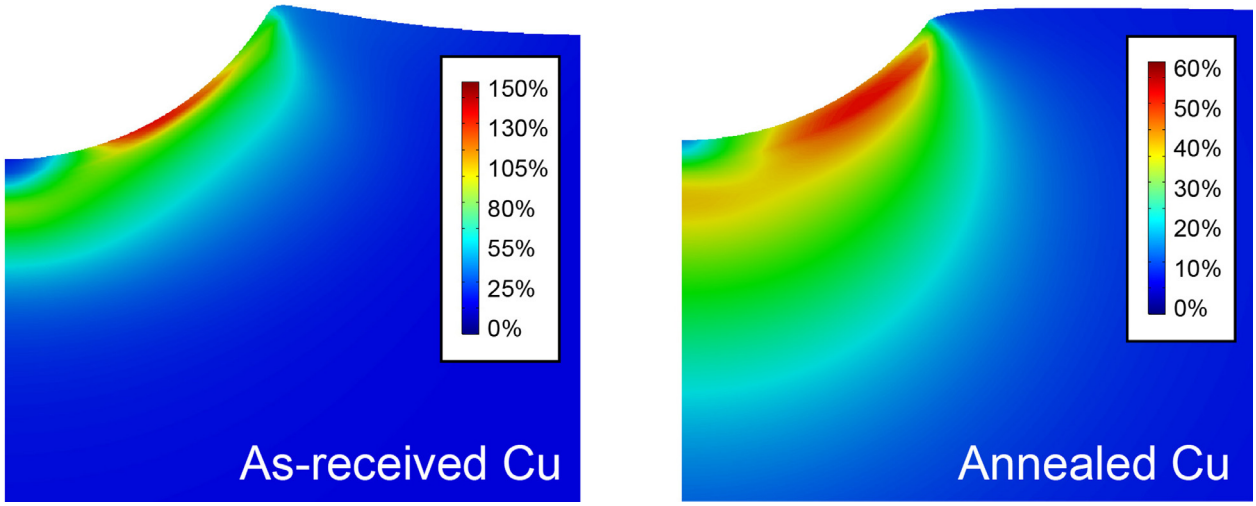

Fig. 6. Predicted plastic strain fields after indentation to a penetration ratio (δ/R) of 40% for as-received and annealed Cu.

suited to this kind of cross-checking, since they are slightly more sensitive to the value of μ than are load-displacement plots.

Such comparisons are shown in Figs. 9 and 10, which relate to both as-received and annealed copper samples. These two materials differ quite substantially in their work hardening characteristics and the comparisons highlight some of the effects of work hardening on the outcomes. Fig. 9 shows load-displacement plots, while Fig. 10 relates to residual indent profiles. Two points are immediately clear from study of these two figures. One is that neglecting friction ($\mu = 0$) leads to significantly different behavior from that predicted for higher values, particularly at relatively high penetration ratios. The other is that, provided μ is at least about 0.2, increasing it further does not make much difference. Moreover, a value of about 0.2 gives good agreement between prediction and experiment for all of the cases studied here. Physically, this may correspond to a value of this order being sufficient to cut interfacial sliding down to a low level for these cases. This is convenient, since it means it's probably acceptable to simply set μ to some suitable value, such as 0.2, which has been used in all of the work presented here. It should be mentioned, however, that higher values could be appropriate under some circumstances, and could affect the behavior. This is an area requiring further work.

6. Conclusions

The following conclusions can be drawn from this work:

(a) The methodology of iterative FEM simulation of the indentation process, with systematic comparison between experimental and predicted outcomes, is basically sound and can in principle be used to infer various mechanical properties. While the present work is focused on plasticity, the approach is also applicable to others, such as creep and superelasticity. It will always be necessary to use a constitutive law of some type, with the objective being to evaluate the parameters in such an equation.

(b) The current work has involved detailed study of experimental and computational aspects of how this methodology can be optimized (for plasticity). It is already recognized that yielding and work-hardening characteristics of most (metallic) materials can be well captured by at least one of the two most commonly-used constitutive laws (Ludwik–Hollomon and Voce), both of which incorporate 3 parameter values. Finding the set of these values that best captures the behavior thus requires searching 3-D parameter space for the location where predicted outcomes of an indentation test match the corresponding experimental outcomes most closely. However, the friction coefficient, μ, during indentation is a further parameter that needs to be included in this optimization. It is shown here that neglect of friction will often introduce a significant error. On the other hand, it appears that, above ~0.15, the effect of it being changed is small. A value of 0.2 has been used in all of the results presented here.

(c) The prime outcome used in the current work has been the load-displacement plot obtained during testing (with a spherical indenter), although the residual indent shape has also been used. An important deduction is that, in order to obtain outcomes that are reasonably sensitive to the work hardening characteristics, relatively deep indenter penetration is required. The penetration ratio, δ/R, should be at least ~25% and preferably ~40%. The other key issue relating to scale is that the sample volume deformed during the test must be large enough to capture the "bulk" response, which usually means that it must contain a sizeable number of grains (perhaps > ~20). This leads to a requirement for relatively large indenters (~ mm range) and loads (~kN range).

(d) The requirement to use relatively large indenters and loads means that many "nano-indenter" systems are not well suited to this methodology. On the other hand, it also means that the experimental procedures are simpler, more transparent and less temperamental than in many fine scale systems. Surface preparation is straightforward - a standard polish to a roughness of a few microns is fine - and there need be few concerns about the effects of oxide layers or surface contamination. Furthermore, obtaining suitable indenters is easy and cheap - cermet spheres in this size range are readily obtainable. Displacement measurement must be accurate, although, again, the relatively coarse scale of the operation means that ultra-high precision is not needed and a standard eddy current gauge or LVDT should be adequate.

(e) It is, however, important that conditions during indentation -

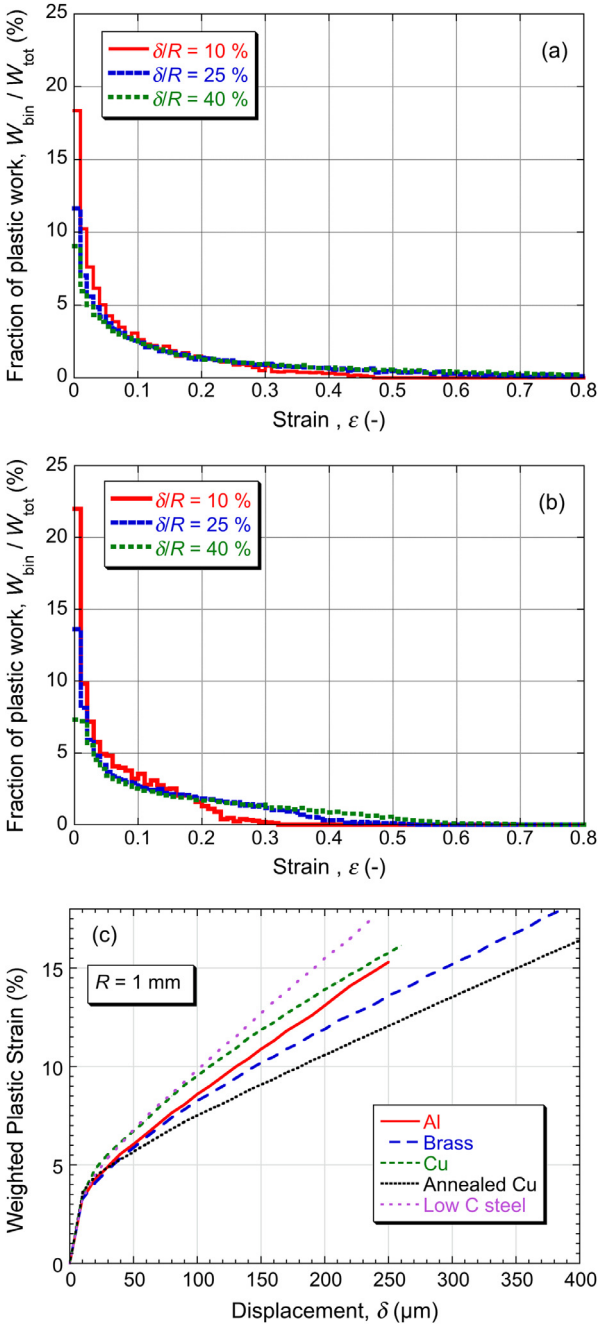

Fig. 7. Fractions of the total work done in different ranges of strain during penetration to 3 different depths into: (a) as-received Cu, (b) annealed Cu, plus (c) average plastic strains for all materials, as a function of the penetration.

particularly the displacement measurements - should closely reflect those simulated in the FEM model. Meshes in software packages for universal use are likely to incorporate only the sample (the exact dimensions of which are not important), perhaps plus the lower half of the indenter. This does mean that a compliance calibration must be applied to indentation load-displacement data. This should include the compliance of the (upper half of the) indenter and its housing, although in many cases this will constitute only a small correction. Other parts of the loading train, however, may in some systems have a relatively large compliance and this must be subtracted. (This also applies to accurate measurement of stress-strain curves by conventional uniaxial testing.)

(f) Information is also presented about the computational approaches and algorithms. Searching of plasticity parameter space is done on

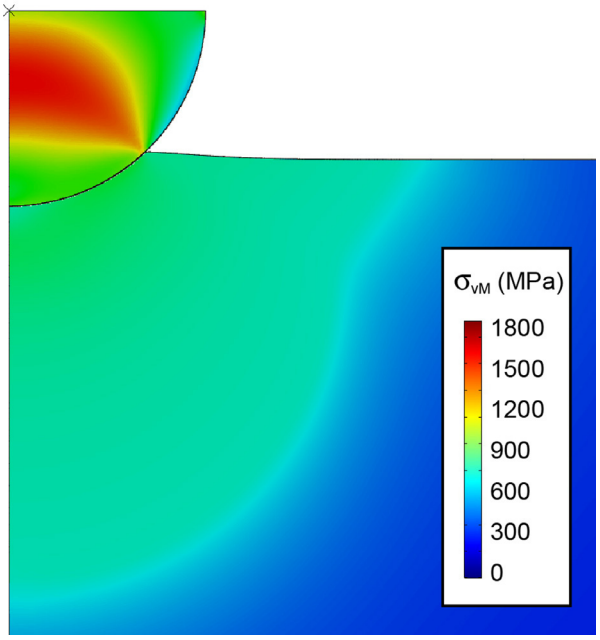

Fig. 8. Predicted distribution of the von Mises stress during indentation of the low C steel, with the penetration ratio having reached 25%. For a 1 mm radius indenter, the applied load at this point would be 3.25 kN.

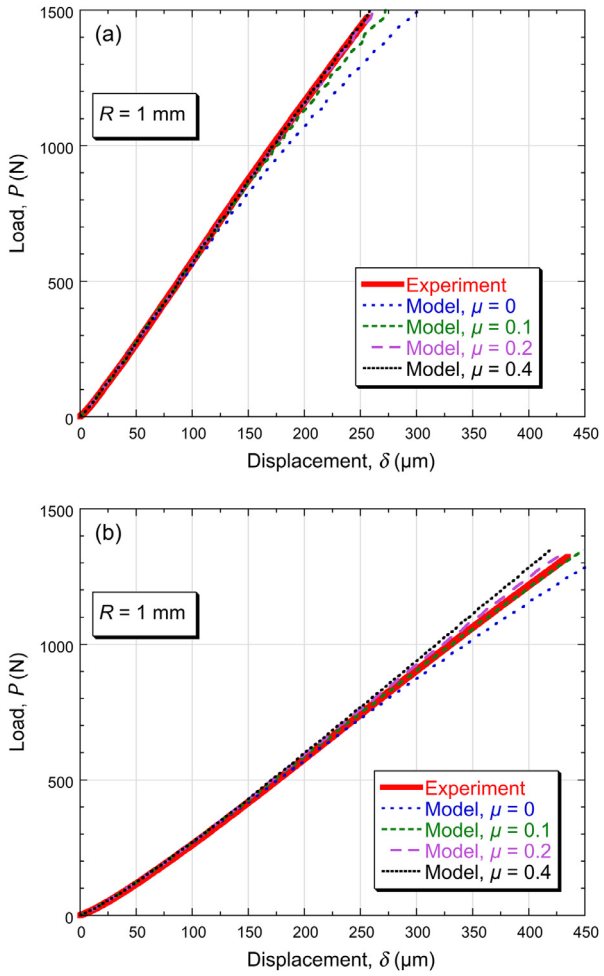

Fig. 9. Comparison between experimental load-displacement plots and corresponding model predictions, obtained using the plasticity parameter values in Table 2 and the μ values shown in the legend, for (a) as-received Cu and (b) annealed Cu.

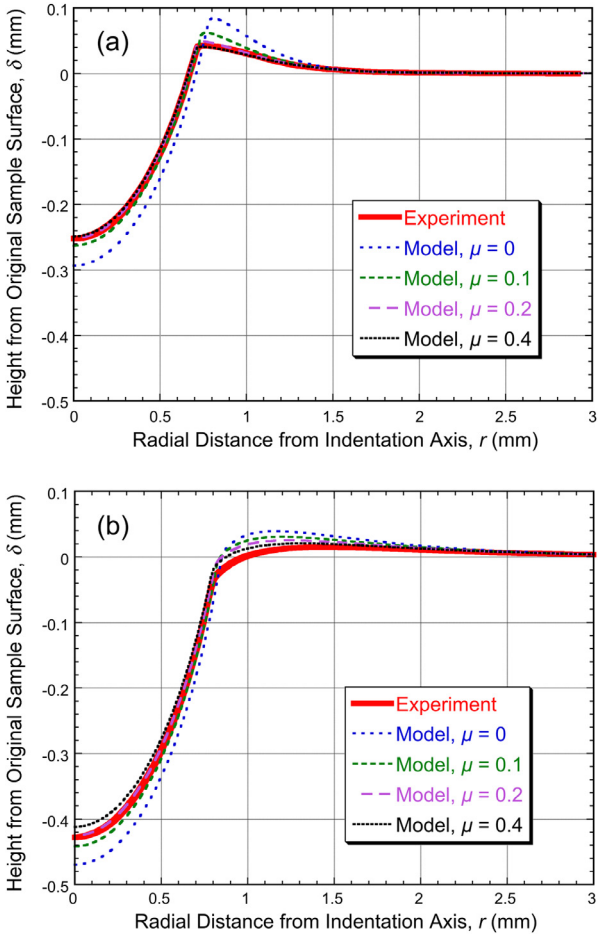

Fig. 10. Comparison between experimental residual indent profiles and corresponding model predictions, obtained using the plasticity parameter values in Table 2 and the μ values shown in the legend, for (a) as-received Cu and (b) annealed Cu.

the basis of a goodness-of-fit parameter and the current work has involved minimizing S (the sum of the squares of the residuals). It is shown that an adaptation of the Nelder–Mead simplex search algorithm is efficient in converging fairly rapidly on the best-fit solution. Computing time requirements will depend on a number of factors, but the procedure is basically a tractable one, with "answers" often obtainable in a matter of minutes (particularly if a "pre-run" matrix of predicted outcomes is available).

(g) These methodologies have been applied to 5 different metals, all in the form of extruded rods. Conventional uniaxial (compressive) testing revealed that they cover a wide range of plasticity characteristics. Some of them exhibit noticeable anisotropy. It is shown that, for all of these materials, the procedure allows these characteristics to be obtained (solely via indentation), with excellent fidelity (± 5% over the complete strain range of interest - i.e. up to ~15–20%). It should be noted, however, that indentation responses are highly-multi-axial and will lead to inferred stress-strain curves that are averages of those obtained uniaxially in different directions. This constitutes a limitation on the information obtainable by indentation, although it does mean that there is no need for any concern about the loading direction relative to the orientation of the sample.

(h) Many of the concepts and approaches described here are already incorporated into usable software packages, which are self-contained in terms of FEM modeling and convergence capabilities. Users are simply required to input appropriate experimental data. In fact, there is already a website available (https://www.sempid.com/about) where a capability of this type is described.

Acknowledgments

This work has been supported by EPSRC (grant RG62695) and also by AWE, as part of an ongoing collaboration aimed at the development of robust and user-friendly tools for the extraction of mechanical property characteristics from instrumented indentation data. The authors are particularly grateful to Giles Aldrich-Smith and Nigel Park, for extensive collaboration and support.

In compliance with current EPSRC recommendations, code and input data for the indentation modeling described here, including meshing and boundary condition specifications, are available at: www.ccg.msm.cam.ac.uk/publications/resources. The code is available at: github.com/rpt26/plasticity_extraction. It includes the convergence algorithm. These files can be downloaded and used in ABAQUS FEM models. Data supplied are for a representative case (with a spherical indenter and radial symmetry).

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.mechmat.2018.06.004.

Appendix: Nelder–Mead convergence algorithm

Once an initial simplex has been created, each iteration comprises the following steps.

1. The values of S are calculated for each vertex and the vertices are ranked, such that $S(\boldsymbol{x}_1) < S(\boldsymbol{x}_2) < \ldots < S(\boldsymbol{x}_{m+1})$. The point to be replaced is $\boldsymbol{x}_{m+1}$. The centroid of the (reduced) simplex is calculated from:

$$\boldsymbol{x}_{\text{cen}} = \frac{1}{m}\sum_{j=1}^{m}\boldsymbol{x}_j \tag{A1}$$

This defines the search direction ($\boldsymbol{x}_{\text{cen}}$ - $\boldsymbol{x}_{m+1}$).

2. **Reflection**: A trial point is established by reflection of $\boldsymbol{x}_{m+1}$ through $\boldsymbol{x}_{\text{cen}}$.

$$\boldsymbol{x}_{\text{ref}} = \boldsymbol{x}_{\text{cen}} + \alpha(\boldsymbol{x}_{\text{cen}} - \boldsymbol{x}_{m+1}) \tag{A2}$$

where α is a scale factor. The value of S is calculated for this point. If $S(\boldsymbol{x}_1) < S(\boldsymbol{x}_{\text{ref}}) < S(\boldsymbol{x}_m)$, so that $\boldsymbol{x}_{\text{ref}}$ is of intermediate quality, then $\boldsymbol{x}_{\text{ref}}$ is accepted, replacing $\boldsymbol{x}_{m+1}$. Otherwise, the algorithm proceeds to step 3.

3. **Expansion**: If $S(\boldsymbol{x}_{\text{ref}}) < S(\boldsymbol{x}_1)$, so that $\boldsymbol{x}_{\text{ref}}$ is the best point yet, this could indicate that the simplex is on an extended downward gradient and an expanded point is trialed

$$\boldsymbol{x}_{\text{exp}} = \boldsymbol{x}_{\text{cen}} + \beta(\boldsymbol{x}_{\text{cen}} - \boldsymbol{x}_{m+1}) \tag{A3}$$

where β is a scale factor ($> \alpha$). The value of S is calculated for this point. If $S(\boldsymbol{x}_{\text{exp}}) < S(\boldsymbol{x}_{\text{ref}})$, then $\boldsymbol{x}_{\text{exp}}$ is accepted, replacing $\boldsymbol{x}_{m+1}$. Otherwise, $\boldsymbol{x}_{\text{ref}}$ is accepted, replacing $\boldsymbol{x}_{m+1}$.

4. **Outside contraction**: If $S(\boldsymbol{x}_m) \leq S(\boldsymbol{x}_{\text{ref}}) < S(\boldsymbol{x}_{m+1})$, so that $\boldsymbol{x}_{\text{ref}}$ is an improvement on $\boldsymbol{x}_{m+1}$, but would become the new worst point, the value of S Is calculated for a point between $\boldsymbol{x}_{\text{ref}}$ and $\boldsymbol{x}_{\text{cen}}$, called the outside contraction point.

$$\boldsymbol{x}_{\text{OC}} = \boldsymbol{x}_{\text{cen}} + \gamma(\boldsymbol{x}_{\text{cen}} - \boldsymbol{x}_{m+1}) \tag{A4}$$

where γ is a scale factor ($< \alpha$). The value of S is calculated for this point. If $S(\boldsymbol{x}_{\text{OC}}) \leq S(\boldsymbol{x}_{\text{ref}})$, then $\boldsymbol{x}_{\text{OC}}$ is accepted, replacing $\boldsymbol{x}_{m+1}$. Otherwise, the algorithm proceeds to step 6.

5. **Inside contraction**: If $S(\boldsymbol{x}_{m+1}) \leq S(\boldsymbol{x}_{\text{ref}})$, so that $\boldsymbol{x}_{\text{ref}}$ is worse than all of the points in the existing simplex, then the value of S Is calculated for a point between $\boldsymbol{x}_{\text{cen}}$ and $\boldsymbol{x}_{m+1}$, called the inside contraction point.

$$\boldsymbol{x}_{\text{IC}} = \boldsymbol{x}_{\text{cen}} - \delta(\boldsymbol{x}_{\text{cen}} - \boldsymbol{x}_{m+1}) \tag{A5}$$

where δ is another scale factor. The value of S is calculated for this point. If $S(\boldsymbol{x}_{\text{IC}}) < S(\boldsymbol{x}_{\text{ref}})$, then $\boldsymbol{x}_{\text{IC}}$ is accepted, replacing $\boldsymbol{x}_{m+1}$. Otherwise, the algorithm proceeds to step 6.

6. **Shrink**: If none of the previous steps are able to improve the simplex, then it is shrunk towards the best vertex. This operation is defined by

$$\boldsymbol{x}'_j = \boldsymbol{x}_j + \delta(\boldsymbol{x}_1 - \boldsymbol{x}_j) \tag{A6}$$

for $2 \leq j \ (m+1)$. The algorithm then starts the next iteration at step 1.

The scale factors (α, β, γ and δ) are often ascribed values of 1, 2, 0.5 and 0.5 respectively, but these can be tuned to cope with particular situations, such as different levels of noise. The Scientific Python implementation allows for these scale factors to be adapted as the algorithm proceeds, as described in Gao and Han (2012).

References

Basu, S., Moseson, A., Barsoum, M.W., 2006. On the determination of spherical nanoindentation stress-strain curves. J. Mater. Res. 21 (10), 2628–2637.

Bobzin, K., Bagcivan, N., Theiss, S., Brugnara, R., Perne, J., 2013. Approach to determine stress strain curves by FEM supported nanoindentation. Materialwiss. Werkstofftech. 44 (6), 571–576.

Bolzon, G., Maier, G., Panico, M., 2004. Material model calibration by indentation, imprint mapping and inverse analysis. Int. J. Solids Struct. 41 (11-12), 2957–2975.

Bouzakis, K., Michailidis, N., 2004. Coating elastic-plastic properties determined by means of nanoindentations and FEM-supported evaluation algorithms. Thin Solid Films 469, 227–232.

Bouzakis, K., Michailidis, N., 2006. An accurate and fast approach for determining materials stress-strain curves by nanoindentation and its FEM-based simulation. Mater. Charact. 56, 147–157.

Bucaille, J.L., Stauss, S., Felder, E., Michler, J., 2003. Determination of plastic properties of metals by instrumented indentation using different sharp indenters. Acta Mater. 51 (6), 1663–1678.

Burley, M., Campbell, J.E., Dean, J., Clyne, T.W., 2018. Johnson-Cook parameter evaluation from ballistic impact data via iterative FEM modelling. Int. J. Impact Eng. 112, 180–192.

Capehart, T.W., Cheng, Y.T., 2003. Determining constitutive models from conical indentation: Sensitivity analysis. J. Mater. Res. 18 (4), 827–832.

Cheng, Y.-T., Cheng, C.-M., 2004. Scaling, dimensional analysis, and indentation measurements. Mater. Sci. Eng. 44 (4-5), 91–149.

Chollacoop, N., Dao, M., Suresh, S., 2003. Depth-sensing instrumented indentation with dual sharp indenters. Acta Mater. 51 (13), 3713–3729.

Dao, M., Chollacoop, N., Van Vliet, K.J., Venkatesh, T.A., Suresh, S., 2001. Computational modeling of the forward and reverse problems in instrumented sharp indentation. Acta Mater. 49, 3899–3918.

Dean, J., Clyne, T.W., 2017. Extraction of plasticity parameters from a single test using a spherical indenter and FEM modelling. Mech. Mater. 105, 112–122.

Dean, J., Wheeler, J.M., Clyne, T.W., 2010. Use of quasi-static nanoindentation data to obtain stress-strain characteristics for metallic materials. Acta Mater. 58, 3613–3623.

Estrin, Y., Mecking, H., 1984. A unified phenomenological description of work hardening and creep based on one-parameter models. Acta Metall. 32 (1), 57–70.

Futakawa, M., Wakui, T., Tanabe, Y., Ioka, I., 2001. Identification of the constitutive equation by the indentation technique using plural indenters with different apex angles. J. Mater. Res. 16, 2283–2292.

Gao, F.C., Han, L.X., 2012. Implementing the Nelder-Mead simplex algorithm with adaptive parameters. Comput. Optim. Appl. 51 (1), 259–277.

Giannakopoulos, A.E., Suresh, S., 1999. Determination of elastoplastic properties by instrumented sharp indentation. Scr. Mater. 40 (10), 1191–1198.

Guelorget, B., Francois, M., Liu, C., Lu, J., 2007. Extracting the plastic properties of metal materials from microindentation tests: experimental comparison of recently published methods. J. Mater. Res. 22, 1512–1519.

Hamada, A.S., Haggag, F.M., Porter, D.A., 2012. Non-destructive determination of the yield strength and flow properties of high-manganese twinning-induced plasticity steel. Mater. Sci. Eng.-a Struct. Mater. Prop. Microstruct. Process. 558, 766–770.

Hausild, P., Materna, A., Nohava, J., 2012. On the identification of stress-strain relation by instrumented indentation with spherical indenter. Mater. Des. 37, 373–378.

Heinrich, C., Waas, A.M., Wineman, A.S., 2009. Determination of material properties using nanoindentation and multiple indenter tips. Int. J. Solids Struct. 46, 364–376.

Herbert, E.G., Pharr, G.M., Oliver, W.C., Lucas, B.N., Hay, J.L., 2001. On the measurement of stress-strain curves by spherical indentation. Thin Solid Films 398, 331–335.

Hollomon, J.H., 1945. Tensile deformation. Trans. Am. Inst. Min. Metall. Eng. 162, 268–290.

Isselin, J., Iost, A., Golek, J., Najjar, D., Bigerelle, M., 2006. Assessment of the constitutive law by inverse methodology: Small punch test and hardness. J. Nucl. Mater. 352 (1-3), 97–106.

Kang, B.S.J., Yao, Z., Barbero, E.J., 2006. Post-yielding stress-strain determination using spherical indentation. Mech. Adv. Mater. Struct. 13 (2), 129–138.

Karthik, V., Visweswaran, P., Bhushan, A., Pawaskar, D.N., Kasiviswanathan, K.V., Jayakumar, T., Raj, B., 2012. Finite element analysis of spherical indentation to study pile-up/sink-in phenomena in steels and experimental validation. Int. J. Mech. Sci. 54 (1), 74–83.

Ma, Z.S., Zhou, Y.C., Long, S.G., Zhong, X.L., Lu, C., 2012. Characterization of stress-strain relationships of elastoplastic materials: An improved method with conical and pyramidal indenters. Mech. Mater. 54, 113–123.

Mecking, H., Kocks, U.F., 1981. Kinetics of flow and strain-hardening. Acta Metall. 29 (11), 1865–1875.

Nelder, J.A., Mead, R., 1965. A simplex method for function minimization. Comput. J. 7 (4), 308–313.

Oliphant, T.E., 2007. Python for scientific computing. Comput. Sci. Eng. 9 (3), 10–20.

Patel, D.K., Kalidindi, S.R., 2016. Correlation of spherical nanoindentation stress-strain curves to simple compression stress-strain curves for elastic-plastic isotropic materials using finite element models. Acta Mater. 112, 295–302.

Pathak, S., Kalidindi, S.R., 2015. Spherical nanoindentation stress-strain curves. Mater. Sci. Eng. R-Reports 91, 1–36.

Pelletier, H., 2006. Predictive model to estimate the stress-strain curves of bulk metals using nanoindentation. Tribol. Int. 39 (7), 593–606.

Riley, K.F., Hobson, M.P., Bence, S.J., 2006. Mathematical Methods for Physics and Engineering. Cambridge University Press.

Sainath, G., Choudhary, B.K., Christopher, J., Samuel, E.I., Mathew, M.D., 2015. Applicability of Voce equation for tensile flow and work hardening behaviour of P92 ferritic steel. Int. J. Press. Vessels Pip. 132, 1–9.

Samuel, K.G., Rodriguez, P., 2005. On power-law type relationships and the Ludwigson explanation for the stress-strain behaviour of AISI 316 stainless steel. J. Mater. Sci. 40 (21), 5727–5731.

Skelton, R.P., Maier, H.J., Christ, H.J., 1997. The Bauschinger effect, masing model and the Ramberg-Osgood relation for cyclic deformation in metals. Mater. Sci. Eng. a-Struct. Mater. Prop. Microstruct. Process. 238 (2), 377–390.

Sun, Y., Zheng, S., Bell, T., Smith, J., 1999. Indenter tip radius and load frame compliance

calibration using nanoindentation loading curves. Phil. Mag. Lett. 79 (9), 649–658.
Taljat, B., Pharr, G.M., 2004. Development of pile-up during spherical indentation of elastic-plastic solids. Int. J. Solids Struct. 41 (14), 3891–3904.
Taljat, B., Zacharia, T., Kosel, F., 1998. New analytical procedure to determine stress-strain curve from spherical indentation data. Int. J. Solids Struct. 35 (33), 4411–4426.
Ullner, C., Reimann, E., Kohlhoff, H., Subaric-Leitis, A., 2010. Effect and measurement of the machine compliance in the macro range of instrumented indentation test. Measurement 43 (2), 216–222.
van der Walt, S., Colbert, S.C., Varoquaux, G., 2011. The NumPy array: a structure for efficient numerical computation. Comput. Sci. Eng. 13 (2), 22–30.
Van Vliet, K.J., Prchlik, L., Smith, J.F., 2011. Direct measurement of indentation frame compliance. J. Mater. Res. 19 (1), 325–331.
Voce, E., 1948. The relationship between stress and strain for homogeneous deformation. J. Inst. Met. 74 (11), 537–562.
Xu, B.X., Chen, X., 2010. "Determining engineering stress-strain curve directly from the load-depth curve of spherical indentation test. J. Mater. Res. 25 (12), 2297–2307.