

Process Parameter Control in Stir Casting of Aluminium Matrix Composites: A Review of Wettability, Particle Distribution and Porosity Mechanisms

Manikandan Ravi

Department of Mechanical Engineering, Saveetha School of Engineering,
Saveetha Institute of Medical and Technical Sciences (SIMATS), Saveetha University,
Chennai, Tamil Nadu, India.
mani131990@gmail.com

Abstract

Stir casting remains the dominant liquid-state route for the manufacture of particulate-reinforced aluminium matrix composites (AMCs), principally because it accommodates large batch sizes, tolerates a wide range of reinforcement chemistries and imposes the lowest capital burden among competing routes. Its persistent limitation is reproducibility: the same nominal composition, processed on different equipment or by different operators, frequently yields divergent microstructures and consequently divergent properties. This review argues that the reproducibility problem is not attributable to any single parameter but to the fact that the governing parameters act on three distinct and partially antagonistic physical mechanisms — wettability at the particle–melt interface, hydrodynamic transport and dispersion within the vortex, and gas entrainment leading to porosity. A parameter setting that improves one mechanism frequently degrades another, and the optimum is therefore a compromise rather than a maximum. The review organises the literature around this three-mechanism framework rather than around parameters considered in isolation. Stirring speed, stirring duration, impeller geometry and immersion depth, melt superheat and processing temperature regime, particle preheating and feed rate, wettability modification through magnesium addition and halide salt fluxing, and atmosphere control are each examined for their mechanistic action and their trade-offs. Characteristic defect morphologies — clustering, settling and pushing segregation, interfacial reaction products and gas porosity — are related back to the parameter combinations that produce them. The optimisation literature, dominated by Taguchi and response surface designs, is assessed critically: the review finds that the prevalence of single-response optimisation and the near-universal omission of melt-temperature-history reporting together account for much of the inter-study inconsistency in reported optima. Five research gaps are identified, of which the absence of standardised distribution-quantification metrics and of validated coupled thermal–hydrodynamic–particle models are judged the most consequential.

Keywords: stir casting; aluminium matrix composites; process parameters; wettability; particle distribution; porosity; metal matrix composites

1. Introduction

Particulate-reinforced aluminium matrix composites occupy a well-defined position in the materials landscape: they offer specific stiffness, wear resistance and elevated-temperature stability superior to monolithic aluminium alloys, while retaining a substantial part of the formability and thermal conductivity that make aluminium attractive in the first place (Ibrahim et al., 1991; Surappa, 2003). Applications in automotive braking, pistons and cylinder liners, and in structural components where weight is at a premium, are established rather than prospective.

The routes available for their manufacture divide broadly into solid-state and liquid-state families. Powder metallurgy offers superior control of reinforcement distribution and avoids the melt-side interfacial reaction problem entirely, but is constrained by component size, tooling cost and the expense of powder handling. Among liquid-state routes, infiltration and spray deposition each address specific niches. Stir casting has nonetheless remained dominant in both industrial and academic practice, and the reason is economic rather than metallurgical: Hashim et al. (1999), in what remains the most cited assessment of the comparison, place stir casting at roughly a third of the cost of competing methods for equivalent output, with the additional advantage of scalability to large batch volumes.

That economic advantage is purchased at the cost of process sensitivity. The literature contains numerous reports in which nominally identical composite systems — the same matrix alloy, the same reinforcement at the same volume fraction — exhibit substantially different microstructures and mechanical responses when produced in different laboratories. The conventional explanation attributes this to insufficiently controlled processing parameters, and the conventional remedy is a parametric optimisation study.

This review takes a different position. The proliferation of parametric optimisation studies has not converged on a consensus parameter set, and the failure of convergence is itself informative. It suggests that the parameters

do not act independently on a single objective, but rather act on multiple competing mechanisms whose relative importance depends on the specific matrix–reinforcement system. A stirring speed that optimises dispersion in one system may, in another with different particle density and wetting characteristics, promote gas entrainment before dispersion is achieved.

The organising contribution of this review is therefore mechanistic rather than parametric. Section 3 establishes three governing mechanisms — interfacial wetting, hydrodynamic dispersion, and gas entrainment — and Sections 4 and 5 examine each processing parameter in terms of its action on all three, with explicit attention to cases in which the actions conflict. Section 6 treats defect formation as the observable consequence of mechanism imbalance. Section 7 assesses the optimisation literature against this framework, and Section 8 identifies the resulting research gaps.

2. Scope and Review Approach

This review addresses the stir casting of aluminium and aluminium alloy matrices reinforced with discontinuous particulates, including ceramic reinforcements (SiC, Al₂O₃, B₄C, TiC, TiB₂), industrial and agricultural waste-derived reinforcements (fly ash, red mud, rice husk ash, bagasse ash), and solid lubricant additions (graphite, MoS₂) in hybrid systems. Continuous-fibre composites, in-situ composite synthesis and non-aluminium matrices are outside the scope except where they illuminate a mechanism under discussion.

The literature is organised by mechanism rather than by chronology or by reinforcement type. Where studies report conflicting optima, the review attempts to identify the source of the conflict rather than to average across it. Quantitative parameter windows reported in the primary literature are cited to their sources and are treated as system-specific rather than as universal recommendations; readers are cautioned that a parameter value optimal for one matrix–reinforcement pair carries no presumption of transferability.

3. Governing Mechanisms

3.1 Wetting at the Particle–Melt Interface

The incorporation of a ceramic particle into a molten aluminium bath is opposed thermodynamically. Molten aluminium exhibits high surface tension, and most ceramic reinforcements present a low-energy surface, so that the contact angle between melt and particle typically exceeds ninety degrees — the non-wetting condition. Under this condition a particle introduced at the melt surface is not drawn into the bath; it is rejected, and remains floating as a dross layer regardless of the vigour of stirring (Hashim et al., 2001).

The situation is aggravated by surface contamination. Ceramic particles as received carry adsorbed moisture and gases, and aluminium melts carry a tenacious surface oxide film. Both interpose a barrier between the melt and the particle surface at which wetting would otherwise occur.

Three intervention strategies appear in the literature. The first is chemical modification of the melt, most commonly through magnesium addition; magnesium is surface-active in aluminium melts and lowers surface tension, and additionally scavenges oxygen from the particle surface. The second is fluxing, in which halide salts — potassium hexafluorotitanate being the most frequently reported — disrupt the oxide film at the interface. The third is modification of the particle rather than the melt, through preheating to desorb surface moisture or through metallic coating to present a wettable surface. These are examined in Section 5.

The mechanistic significance of wetting extends beyond incorporation. A poorly wetted particle, even if mechanically entrained by the vortex, forms a weak interface with the solidified matrix, and that interface functions as a crack initiation site under load. Wetting therefore governs both processability and the mechanical property outcome, and this dual role explains why wettability treatment frequently produces property improvements disproportionate to the modest compositional change involved.

3.2 Hydrodynamic Transport and Dispersion

Once incorporated, particles must be distributed. The vortex generated by an impeller performs two distinct functions that are worth separating conceptually: it draws surface-floating particles into the bulk melt through the vortex core, and it disperses agglomerates through shear.

Naher et al. (2003, 2007), in combined computational and experimental work, established that the flow field in a stirred melt is markedly inhomogeneous, with shear concentrated near the impeller and comparatively quiescent zones near the crucible wall and base. Particles resident in the quiescent zones experience little dispersive shear, which explains the persistence of clustering at nominally adequate stirring speeds and durations.

Two transport processes oppose uniform distribution. Density-driven separation — settling for reinforcements denser than the melt, flotation for lighter ones such as fly ash cenospheres — acts throughout the period between stirring cessation and solidification. The relevant comparison is between the characteristic separation time and the solidification time, which places pouring temperature and mould thermal conditions within the distribution problem rather than outside it. Separately, the advancing solidification front interacts with

suspended particles: below a critical front velocity the particle is pushed ahead of the interface and accumulates in the last-freezing regions, producing interdendritic segregation, while above it the particle is engulfed and remains where the melt held it. Slow cooling therefore promotes segregation independently of how well the melt was stirred.

3.3 Gas Entrainment and Porosity

The third mechanism is the one most frequently neglected in parametric studies. A vortex of sufficient depth entrains atmosphere at its core; the entrained gas is dispersed as bubbles by the same shear field that disperses particle agglomerates, and a substantial fraction is retained through solidification as porosity. Hydrogen dissolution from atmospheric moisture, promoted by the increased melt surface area that stirring produces, contributes further.

Porosity in stir cast composites therefore has three superimposed sources: mechanically entrained air, dissolved hydrogen precipitated on solidification, and gas adsorbed on particle surfaces and released into the melt. Each responds to different parameters, which is a principal reason that porosity correlations derived from single-parameter studies transfer poorly.

The mechanistic conflict central to this review can now be stated precisely. Increased stirring speed and duration improve dispersion (Section 3.2) and assist incorporation of marginally wetting particles (Section 3.1), while simultaneously increasing gas entrainment (Section 3.3). No parameter setting maximises all three mechanisms. The task of process design is the location of an acceptable compromise, and the position of that compromise depends on which mechanism is limiting for the specific system — which in turn depends on particle density, size, surface chemistry and volume fraction.

4. Hydrodynamic Parameters

4.1 Stirring Speed

Stirring speed is the most extensively studied parameter and the one for which reported optima diverge most widely. The mechanistic picture accounts for the divergence. At low speeds the vortex is shallow and the shear rate low; incorporation is incomplete and agglomerates survive. As speed increases, both incorporation and dispersion improve. Beyond a threshold, vortex depth becomes sufficient for sustained atmospheric entrainment, and porosity rises. Above this the marginal dispersion benefit is small, since agglomerates surviving at high shear are typically bound by sintering or by interlocking rather than by weak adhesion, and further speed increase does not break them.

Prabu et al. (2006) examined stirring speed and duration jointly for an Al–SiC system and reported that distribution homogeneity improved with both, with speed exhibiting the stronger effect and the two showing interaction — an important finding, since it indicates that speed and duration are not interchangeable in the manner that a fixed “total shear” assumption would imply.

The practical inference is that the reported optima in the literature — which cluster loosely in the mid-hundreds of revolutions per minute for laboratory-scale crucibles — are not transferable to different geometries. The mechanistically relevant quantities are the shear rate distribution and the vortex depth relative to melt depth, neither of which is fixed by rotational speed alone. Reporting practice in the literature would be materially improved by the inclusion of impeller diameter and melt depth alongside rotational speed, permitting at least an approximate dimensionless comparison.

4.2 Stirring Duration

Duration exhibits diminishing returns with a defined penalty. The dispersive benefit is concentrated in the early period, as agglomerates are broken and particles distributed; thereafter the incremental homogeneity gain is small. The penalties accumulate linearly or worse: extended residence at temperature promotes interfacial reaction (Section 6.3), extended exposure of the stirred surface promotes hydrogen pickup, and extended processing incurs melt temperature drift unless actively controlled.

In hybrid systems containing graphite or other low-density solid lubricants, prolonged stirring is additionally counterproductive, since the density differential drives flotation of the lighter phase during any interval of reduced agitation.

4.3 Impeller Geometry, Blade Angle and Immersion Depth

Impeller design receives markedly less attention than speed and duration, despite governing the flow field that speed merely scales. Ramanathan et al. (2019), in a review addressed specifically to furnace and stirrer design, note the diversity of impeller configurations in use — flat blade, pitched blade, turbine and helical — and the near-absence of systematic comparison between them under otherwise identical conditions.

The mechanistic considerations are straightforward even where the empirical literature is thin. Blade pitch determines the ratio of axial to radial pumping: axial flow circulates melt vertically and is effective against settling, while radial flow generates higher local shear and is more effective against agglomeration. Immersion depth positions the high-shear zone within the melt; an impeller placed too near the surface produces a deep vortex with strong entrainment and leaves the lower melt poorly agitated, while one placed too near the base leaves surface-floating particles unincorporated. The frequently cited practice of positioning the impeller at approximately one-third of melt depth from the base represents an empirical compromise between these failure modes rather than an optimised value.

Blade number and diameter ratio similarly affect the flow field. This review identifies impeller design as the parameter with the largest gap between mechanistic importance and research attention, and Section 8 returns to the point.

Table 1. Hydrodynamic parameters: mechanistic action and trade-offs

Parameter	Effect on wetting / incorporation	Effect on dispersion	Effect on porosity	Net trade-off
Stirring speed ↑	Improves — assists mechanical entrainment	Improves up to a threshold; marginal thereafter	Degrades — vortex entrainment increases	Optimum below the dispersion plateau
Stirring duration ↑	Neutral after incorporation	Improves early, then diminishing	Degrades — cumulative gas pickup	Short of the plateau; penalties are cumulative
Radial (flat-blade) impeller	Neutral	Higher local shear; better de-agglomeration	Higher surface disturbance	Favoured where clustering dominates
Axial (pitched-blade) impeller	Neutral	Better vertical circulation; resists settling	Lower entrainment	Favoured where density separation dominates
Immersion depth ↑ (deeper)	Degrades — surface particles unincorporated	Lower melt better served	Reduced entrainment	Compromise position required

5. Thermal and Chemical Parameters

5.1 Melt Temperature and Processing Regime

Melt temperature acts on all three mechanisms simultaneously, which makes it the most consequential single parameter and the least satisfactorily reported. Higher superheat lowers viscosity, improving particle transport and reducing the torque required for a given shear rate; it also improves wetting kinetics. It simultaneously accelerates interfacial reaction, increases hydrogen solubility, and — by lengthening the interval to solidification — extends the period available for density-driven separation.

Three processing regimes are distinguished in the literature. Fully liquid processing above the liquidus offers the lowest viscosity and the simplest handling, at the cost of maximal separation and reaction. Semi-solid or compocasting processing within the freezing range exploits the elevated apparent viscosity of the partially solid slurry to suppress settling and to enhance mechanical de-agglomeration through interparticle collision with primary solid; the penalty is greatly increased stirring torque and restricted pourability. Two-step processing, in which the melt is cooled into the semi-solid range for particle addition and stirring before reheating for pouring, attempts to capture both advantages and is reported favourably for nanoparticle and fine-particle systems where agglomeration is severe (Sajjadi et al., 2011).

A reporting deficiency deserves emphasis. Studies frequently specify a single melt temperature, typically the value at particle addition, and omit the subsequent thermal history — the drift during stirring, the temperature at pouring, and the mould preheat condition. Since separation and pushing segregation are governed by the interval between stirring cessation and solidification completion, this omission removes information necessary to interpret the reported microstructure. This review identifies incomplete thermal-history reporting as a probable contributor to inter-study inconsistency, and recommends that the melt temperature at particle addition, at stirring completion and at pouring, together with mould preheat, be reported as a minimum.

5.2 Particle Preheating and Feed Rate

Preheating of reinforcement particles before addition serves three purposes: desorption of surface moisture and adsorbed gases, reduction of the thermal shock imposed on the melt at addition, and — for certain oxide-forming systems — development of a thin surface oxide that improves wettability relative to the as-received surface. The

moisture argument is the strongest, since adsorbed water introduced into an aluminium melt dissociates and contributes hydrogen directly to the porosity budget.

Feed rate is comparatively neglected. Bulk addition produces local particle concentrations at which agglomeration is essentially assured and a local melt temperature depression that raises viscosity precisely where dispersion is required. Continuous or incremental addition at a controlled rate, with stirring maintained throughout, is the mechanistically preferable practice. Quantitative guidance on feed rate is largely absent from the literature.

5.3 Wettability Modification

Magnesium addition is the most widely adopted intervention. Its action is dual: as a surface-active element it depresses the melt surface tension, and as a strong reducing agent it scavenges oxygen at the particle surface, disrupting the barrier oxide. Reported additions are commonly in the low single-digit weight percent range. The intervention is not without cost — magnesium is volatile at melt temperatures, contributes to dross formation, and in excess alters the matrix alloy chemistry and precipitation response.

Halide salt fluxing, principally with potassium hexafluorotitanate, acts through disruption of the oxide film and is reported to be effective at small additions. The residue chemistry and its effect on subsequent corrosion behaviour are less thoroughly characterised.

Particle-side modification through electroless metallic coating (copper and nickel are the most reported) presents a metallic, readily wetted surface to the melt and is highly effective, but the process cost is substantial and largely negates the economic rationale for selecting stir casting in the first place. Its use is consequently confined to systems in which no melt-side intervention proves adequate.

5.4 Atmosphere Control and Degassing

Processing under inert cover gas addresses the gas-entrainment mechanism at its source, since entrained inert gas contributes no dissolved hydrogen and the melt surface is protected from atmospheric moisture. The cost is equipment complexity. Post-stirring degassing, whether by inert gas purging or by hexachloroethane tablet, is widely practised and effective against dissolved hydrogen, but has limited effect on mechanically entrained bubbles already dispersed through a viscous particle-laden melt — a distinction not always observed in the literature. Vacuum processing addresses both, at correspondingly greater cost.

Post-solidification interventions form a separate category. Squeeze casting applies pressure during solidification and is highly effective against gas porosity, with the additional benefit of improved interfacial contact. Ultrasonic treatment of the melt exploits cavitation to break agglomerates at a scale inaccessible to mechanical shear and is increasingly reported for nanoparticle-reinforced systems. Secondary processing — extrusion, rolling or friction stir processing — closes porosity and redistributes clusters, and is the most reliable route to homogeneity where component geometry permits it. It is worth stating plainly that a substantial fraction of the property improvements attributed in the literature to stir casting parameter optimisation could be obtained more reliably through secondary processing.

6. Defect Morphologies and Their Parametric Origins

6.1 Clustering and Agglomeration

Clustering originates in three distinct processes with different parametric remedies, and the failure to distinguish them accounts for some inconsistency in the literature. As-received agglomerates, bound by van der Waals forces or by prior sintering, require dispersive shear for breakup and respond to stirring speed and impeller design. Addition-induced clusters, formed by locally excessive particle concentration at the feed point, respond to feed rate and addition method. Solidification-induced clusters, formed by particle pushing into interdendritic regions, respond to cooling rate and are unaffected by stirring parameters. A study that improves stirring while leaving cooling rate uncontrolled may therefore report negligible improvement in a system where the third mechanism dominates.

6.2 Porosity

Porosity in stir cast AMCs is typically the dominant limitation on mechanical property realisation, since pores act as stress concentrators and crack initiation sites and negate a substantial part of the strengthening contribution the reinforcement was introduced to provide. Its parametric origins are distributed across stirring intensity and duration (entrainment), atmosphere and particle preheating (hydrogen), particle volume fraction and morphology (surface area and interstitial trapping), and pouring and solidification conditions (escape opportunity). The multiplicity of origins is the reason porosity correlations derived from single-parameter studies generalise poorly.

6.3 Interfacial Reaction Products

The Al–SiC system provides the canonical case. Prolonged contact between molten aluminium and silicon carbide above a threshold temperature produces aluminium carbide at the interface, a brittle phase which additionally is hygroscopic and thus degrades corrosion performance. The reaction is time- and temperature-dependent, which places an upper bound on the combination of melt temperature and stirring duration independent of any hydrodynamic consideration. Mitigation strategies include silicon enrichment of the matrix to shift the reaction equilibrium, particle pre-oxidation to interpose a silica layer, and simple limitation of the temperature–time product. Analogous constraints apply to other reinforcement systems and should be established for the specific pair before parameter optimisation is attempted, since an optimum located outside the reaction-safe window is not usable.

6.4 Density-Driven Segregation

Macroscopic gradients in reinforcement concentration along the casting axis arise from settling or flotation during the interval between stirring cessation and solidification. Lightweight reinforcements — fly ash cenospheres in particular — present the flotation case and are frequently reported to concentrate in the upper region of the casting. The controlling variables are the duration of the quiescent interval and the melt viscosity during it, which returns the discussion to pouring temperature and mould thermal conditions. Rapid transfer to a preheated mould with controlled extraction is the practical remedy; semi-solid processing addresses the same problem through viscosity.

Table 2. Defect morphologies and their principal parametric origins

Defect	Dominant mechanism	Principal controlling parameters	Secondary remedy
As-received agglomerates	Insufficient dispersive shear	Stirring speed; impeller geometry; duration	Ultrasonic treatment
Addition-site clusters	Local concentration and chilling at feed point	Feed rate; particle preheat; addition method	Incremental addition
Interdendritic clusters	Particle pushing by solidification front	Cooling rate; mould preheat; pouring temperature	Secondary processing
Gas porosity	Vortex entrainment and hydrogen dissolution	Stirring speed and duration; atmosphere; preheat	Degassing; squeeze casting
Interfacial reaction phase	Temperature–time exposure	Melt temperature; stirring duration	Matrix Si enrichment; particle coating
Axial segregation	Density-driven separation before solidification	Pouring temperature; transfer time; mould condition	Semi-solid processing

7. Optimisation Methodologies: A Critical Assessment

The parametric optimisation literature is dominated by Taguchi orthogonal array designs and by response surface methodology, with analysis of variance used to apportion contribution among factors. The attraction is clear: both permit a defensible conclusion from a modest number of melts, which matters where each experimental run consumes substantial material and furnace time.

Four limitations recur, and they are worth stating because they bear on the interpretation of the reported optima.

First, single-response optimisation predominates. A design that maximises hardness, or minimises wear rate, treats the remaining responses as unconstrained. Given the mechanistic conflicts established in Section 3 — in which dispersion and porosity respond oppositely to stirring intensity — a single-response optimum is very likely to be located at an unfavourable position on the neglected response. Multi-response approaches, whether through desirability functions, grey relational analysis or Pareto formulations, are available and comparatively little used.

Second, factor ranges are typically narrow and centred on values inherited from prior studies. Where the range does not span the mechanistic transition — for stirring speed, the onset of significant vortex entrainment — the fitted response surface is locally valid and the identified optimum is an artefact of the range selected.

Third, screening designs assuming factor independence are applied to a system in which interaction is mechanistically expected. Temperature and duration interact through the reaction kinetics of Section 6.3; speed and duration interact as Prabu et al. (2006) observed directly. Main-effect contribution percentages reported from designs incapable of resolving these interactions should be treated with corresponding caution.

Fourth, and most consequentially for the field, response definitions are not standardised. “Uniform distribution” is variously assessed by qualitative micrograph inspection, by area-fraction variance across fields, by

nearest-neighbour distance statistics or by quadrat methods, and results obtained under these definitions are not comparable. Machine learning approaches now entering the field inherit this problem rather than solving it: a model trained on aggregated literature data acquires the inconsistency of the underlying response definitions, and its apparent predictive accuracy is a statement about the training corpus rather than about the process.

8. Research Gaps and Recommended Directions

- **Standardised distribution metrics.** The absence of an agreed quantitative measure of particle distribution homogeneity is the single greatest obstacle to cumulative progress in this field, since it prevents comparison across studies and corrupts any attempt at meta-analysis or data-driven modelling. Adoption of a defined statistical measure, reported with the field size and magnification at which it was obtained, would be a low-cost intervention with disproportionate benefit.
- **Validated coupled process models.** Existing computational work addresses the melt flow field with reasonable fidelity but generally omits the coupling to particle transport, to thermal history and to solidification-front interaction. A model coupling all four, validated against quantified distribution measurements, would substantially reduce dependence on empirical parameter search and would permit geometry-independent scaling.
- **Impeller design.** The disparity between the mechanistic importance of impeller geometry and the volume of research addressing it is the clearest opportunity in the field. Systematic comparison of blade configurations under otherwise identical conditions, reported with dimensionless flow characterisation, is largely absent.
- **Scale-up characterisation.** Nearly all reported work is at laboratory crucible scale. The shear rate distribution, thermal gradient and solidification time of an industrial-scale melt differ substantially, and the transferability of laboratory optima is essentially uncharacterised.
- **Waste-derived reinforcement variability.** Fly ash, red mud and agricultural ash reinforcements are attractive on cost and sustainability grounds, but their composition, particle size distribution and surface chemistry vary with source and with batch in a manner that ceramic reinforcements do not. Process windows established for one source may not transfer to another, and the sensitivity of the process to this variability has not been systematically established. This is a prerequisite for industrial adoption of waste-reinforced composites rather than an incremental refinement.

9. Conclusions

This review has organised the stir casting parameter literature around three governing mechanisms rather than around parameters considered individually. The principal conclusions are as follows.

- Wetting, dispersion and gas entrainment respond differently and in part oppositely to the principal process parameters, and no parameter set optimises all three. Process design is the location of a compromise whose position is system-specific.
- The wide divergence of reported optima across the literature is consistent with this mechanistic picture and does not require explanation in terms of experimental error. Which mechanism is limiting depends on the particle density, size, volume fraction and surface chemistry of the specific system.
- Melt temperature is the most consequential single parameter, acting on all three mechanisms, and is also the most incompletely reported. Full thermal-history reporting should be regarded as a minimum standard.
- Impeller geometry governs the flow field that rotational speed merely scales, and is substantially under-researched relative to its mechanistic importance.
- The optimisation literature is limited by single-response formulations, narrow factor ranges, unresolved interactions and non-standardised response definitions. Machine learning approaches applied to this corpus inherit rather than resolve these limitations.
- Secondary processing addresses porosity and clustering more reliably than parameter optimisation does, and its availability should inform the decision of how much optimisation effort is warranted.

References

- Hashim, J., Looney, L., Hashmi, M.S.J., 1999. Metal matrix composites: production by the stir casting method. *Journal of Materials Processing Technology* 92–93, 1–7.
- Hashim, J., Looney, L., Hashmi, M.S.J., 2001. The wettability of SiC particles by molten aluminium alloy. *Journal of Materials Processing Technology* 119 (1–3), 324–328.
- Hashim, J., Looney, L., Hashmi, M.S.J., 2002. Particle distribution in cast metal matrix composites — Part I and Part II. *Journal of Materials Processing Technology* 123 (2), 251–257; 258–263.
- Ibrahim, I.A., Mohamed, F.A., Lavernia, E.J., 1991. Particulate reinforced metal matrix composites — a review. *Journal of Materials Science* 26, 1137–1156.

- Kaczmar, J.W., Pietrzak, K., Włosiński, W., 2000. The production and application of metal matrix composite materials. *Journal of Materials Processing Technology* 106 (1–3), 58–67.
- Naher, S., Brabazon, D., Looney, L., 2003. Simulation of the stir casting process. *Journal of Materials Processing Technology* 143–144, 567–571.
- Naher, S., Brabazon, D., Looney, L., 2007. Computational and experimental analysis of particulate distribution during Al–SiC MMC fabrication. *Composites Part A: Applied Science and Manufacturing* 38 (3), 719–729.
- Prabu, S.B., Karunamoorthy, L., Kathiresan, S., Mohan, B., 2006. Influence of stirring speed and stirring time on distribution of particles in cast metal matrix composite. *Journal of Materials Processing Technology* 171 (2), 268–273.
- Ramanathan, A., Krishnan, P.K., Muraliraja, R., 2019. A review on the production of metal matrix composites through stir casting — furnace design, properties, challenges, and research opportunities. *Journal of Manufacturing Processes* 42, 213–245.
- Rohatgi, P.K., 1994. Low-cost, fly-ash-containing aluminium-matrix composites. *JOM* 46 (11), 55–59.
- Sajjadi, S.A., Ezatpour, H.R., Beygi, H., 2011. Microstructure and mechanical properties of Al–Al₂O₃ micro and nano composites fabricated by stir casting. *Materials Science and Engineering: A* 528 (29–30), 8765–8771.
- Singh, J., Chauhan, A., 2016. Overview of wear performance of aluminium matrix composites reinforced with ceramic materials under the influence of controllable variables. *Ceramics International* 42 (1), 56–81.
- Surappa, M.K., 2003. Aluminium matrix composites: challenges and opportunities. *Sadhana* 28 (1–2), 319–334.
- Zhou, W., Xu, Z.M., 1997. Casting of SiC reinforced metal matrix composites. *Journal of Materials Processing Technology* 63 (1–3), 358–363.