Mitigating Metaphors:
A *Comprehensible* Guide to Recent Nature-Inspired Algorithms

Michael A. Lones*

Abstract

In recent years, there has been an explosion of new metaheuristic algorithms that explore different sources of inspiration within the biological and natural worlds. A particular issue with this approach is the tendency for authors to use terminology that is derived from the domain of inspiration, rather than the broader domains of metaheuristics and optimisation. This, in turn, makes it difficult to both comprehend these algorithms and understand their relationships to other metaheuristics. This guide attempts to address this issue, at least to some extent, by providing descriptions of the 32 most cited nature-inspired algorithms published in the last 20 years using standard metaheuristic terms. It also discusses commonalities between these algorithms and more classical nature-inspired metaheuristics such as evolutionary algorithms and particle swarm optimisation.

1 Introduction

This paper is intended to be an objective guide to the most popular nature-inspired optimisation algorithms published since the year 2000, measured by citation count. It is not the first paper to review this area (Fister Jr et al., 2013; Xing and Gao, 2016; Yang, 2010), but it is the first to present these algorithms in terms that will be familiar to the broader optimisation, metaheuristics, evolutionary computation, and swarm computing communities. Unlike some previous reviews, it does not aim to advocate for this area of research or provide support for the idea of designing algorithms based upon observations of natural systems. It only aims to report and summarise what already exists in more accessible terms.

In particular, the aim of this paper is not to criticise these approaches; other authors have already done this for nature-inspired metaheuristics in general (Sörensen, 2015) and for specific nature-inspired algorithms (Črepinšek et al., 2012; Weyland, 2015). However, it is important to be aware of one point of criticism that was raised by Sörensen. This

---

*School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, UK

is the tendency for authors to present their algorithm from the perspective of, and using the terminology of, the domain of inspiration. Often nature-inspired algorithm papers begin with an initial review of a natural domain, then abstract this into a model of the domain, and this leads to an algorithmic description that contains terms from the domain. In many cases, this includes the introduction of new terms to describe well-established concepts from metaheuristics and optimisation. The consequence of this (and something I can personally attest to, having read many these papers) is that it can take considerable time and effort to understand how these algorithms work, even if the reader has a background in metaheuristics.

Nevertheless, these algorithms have gained a significant uptake. This can be seen in their citation counts: the 32 algorithms reviewed in this paper each have more than 200 citations; a third of them have more than 1000 citations. Given that most computer science papers achieve only a handful of citations per year, this is quite an achievement for a group of papers with an average age of around 9 years. However, this combination of high uptake and opaque descriptions has led to fragmentation between the nature-inspired optimisation community and the wider metaheuristics, evolutionary computation, and swarm computing communities. To raise an observation that should be familiar to these communities: a certain amount of diversification is generally a good thing, but diversification without intensification tends to be ineffective. Applying this observation to the design of optimisation algorithms suggests that focusing on variants of a single nature-inspired algorithm is likely to be a sub-optimal approach and a potential waste of time and effort. This, in turn, suggests a need to tie back together these different threads of search. This guide aims to contribute towards this goal.

The main contribution is an annotated glossary of nature-inspired algorithms, with the main characteristics of each of these algorithms described using terminology that will hopefully be familiar to anyone from a metaheuristics or optimisation background. The intent is for this to be used as a resource where someone can look up a particular algorithm and quickly gain an understanding of its main characteristics. Following the listing, there is a short discussion. This highlights areas where algorithms have strong overlap in terms of the mechanisms they use to explore search spaces, both to one another and to more established algorithms. It also highlights areas of novelty.

2 Algorithms from A–Z

Well over a hundred nature-inspired algorithms have been published since 2000; for instance, the review book by Xing and Gao (2016) names 134 of these. It would be a challenging task to read through and understand all these algorithms. Perhaps in reflection of this, previous reviews have generally described these algorithms using their original authors' words, or have focused on the sources of inspiration rather than their metaheuristic mechanisms. By comparison, this guide aims to be comprehensi**ble** rather than comprehensi**ve**. Consequently, it focuses on the more popular of these algorithms. Popularity is measured in terms of citation count; this is not, of course, a robust measure of uptake, but it is probably indicative. To bring the list of algorithms down to a
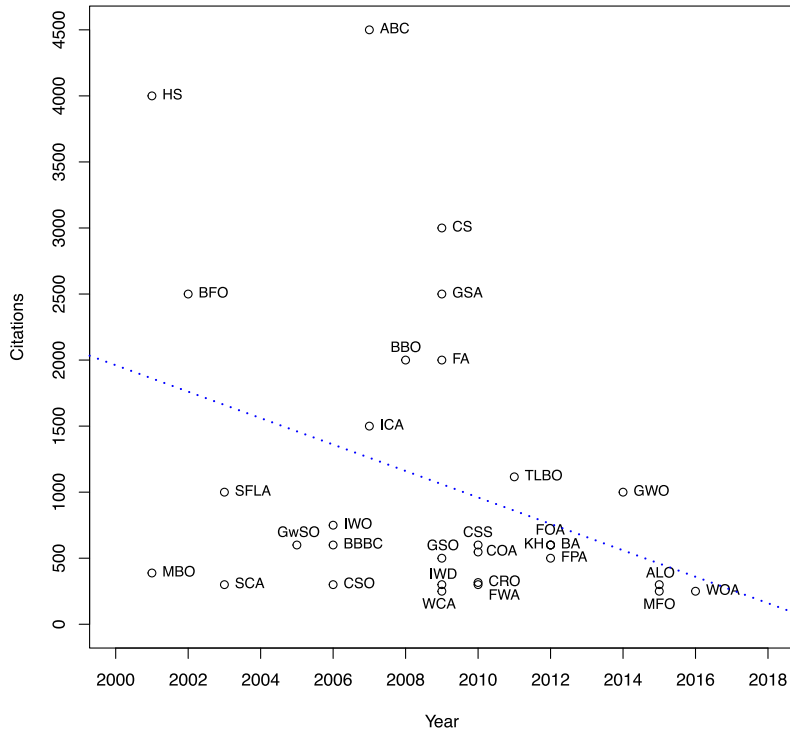
Figure 1: Citation counts of seminal papers, labelled with each algorithm's acronym. See Section 2 for full names and descriptions.

manageable level, this review only covers those which have at least 200 citations, as measured by Google Scholar[1]. By comparison, the seminal genetic algorithm (GA) work has ~60000 citations, particle swarm optimization (PSO) has ~50000 citations, ant colony optimization (ACO) has ~10000 citations, and evolution strategies (ES) have ~5000. It is notable that a number of algorithms in the list have citation counts approaching that of ES and collectively they have ~30000 citations, roughly halfway between the citation counts of ACO and PSO. So, even taking into account the limitations of citation counts, they are clearly having an impact within the scientific record, and this alone should justify efforts to document and understand them. Fig. 1 plots the approximate number of citations against the year that an algorithm's seminal paper was published. It can be seen from the trend line that the average citation count per year is ~100.

The glossary below gives an overview of these 32 algorithms. A brief description is given for each one, written using standard metaheuristic terms, including those which have been standardised by well-known population-based metaheuristics such as evolutionary algorithms (EAs) and PSO. Unless indicated otherwise, it is assumed that each

---

[1]Citations counts were collected in August 2018.

algorithm is a population-based optimiser which updates the population synchronously over a period of iterations and begins with a population that is uniformly sampled from the search space. The phrase "PSO-style algorithm" is used as shorthand to refer to those which use a vector space and have operators that direct search towards other search points within this space. It implies neither that they use all the mechanisms of PSO nor that they do not use mechanisms from other common metaheuristics. The term *particle* is used both for PSO-style algorithms and other algorithms where there is a notion of persistent state attached to a search point. Little or no reference is made to an algorithm's source of inspiration from nature, unless this is required to understand the algorithm. The description is intended to be sufficient to indicate the general characteristics of the algorithm, and to allow the reader to draw out similarities with other algorithms. The descriptions are not intended to be exhaustive, and the reader is encouraged to consult the primary sources if they wish to know more.

Ant Lion Optimizer, ALO, Mirjalili, 2015b, >300 citations

PSO-style algorithm. At each iteration, particles with relatively low fitness are moved to hyper-spheres centred around target particles with relatively high fitness and carry out random walks within these regions of the search space. Target particles are selected in a fitness-proportionate manner so that search points are reallocated to the currently most promising regions. The radius of the hyper-spheres decrease over time to intensify search.

Artificial Bee Colony Algorithm, ABC, Karaboga, 2005, >4500 citations

PSO-style algorithm. For each particle, local moves are made in the direction of another randomly selected particle. Only improving moves are accepted. The number of local moves made from a particular particle is determined by their relative fitness within the population. Particles which have not made progress over a certain number of moves are moved to another, randomly sampled, position. Move sizes are probabilistic and are progressively reduced over time.

Bacterial Foraging Optimization, BFO, Passino, 2002, >2500 citations

At each iteration, each particle carries out a biased random walk: if a move increases fitness, then a further move is made in the same direction; otherwise, a random change in direction occurs. Fitness values are adjusted by a crowding term, whose effect is to draw the random walks towards one another. After each iteration, relatively unfit particles are removed and replaced with copies of relatively fit particles. Occasionally, a random subset of particles are restarted.

Bat Algorithm, BA, Yang and Gandomi, 2012, >600 citations

PSO-style algorithm. Particles are directed towards the population best at different velocities. The magnitude of each particle's velocity is varied randomly at each iteration. There is also a probability of changing the particle's position to a random position near the global best, with the likelihood of doing this decreasing each time the particle makes an improvement. The new position is then accepted probabilistically, with a likelihood that increases each time the particle makes an improvement.

Bees Algorithm, BeA, Pham et al., 2006, >1000 citations

At each iteration, solutions are randomly sampled within a fixed radius of the fittest members of the existing population. The radius reduces progressively over time and only improved solutions are accepted. Random restarts are used to maintain diversity. Similar to an ES with a time-dependent mutation strategy.

Big Bang-Big Crunch, BB-BC, Erol and Eksin, 2006, >600 citations

---

At the start of each iteration, a point representing the fitness-weighted average of the previous population (its "centre of mass") is calculated and a new population is created by sampling from a normal distribution centred around this point. The width of the distribution is reduced at each iteration, intensifying search. Can be seen as an estimation of distribution algorithm (EDA) with a very simple generative model.

Biogeography-Based Optimizer, BBO, Simon, 2008, ~2000 citations

---

At each iteration, for each solution in the population, components are replaced by copying them from other solutions; the likelihood of this is proportional to the solution's fitness, and the likelihood of choosing another solution as a source is proportional to its fitness. Elitism prevents the fittest solutions from being modified. Mutation occurs with probability inversely proportional to fitness. Similarities to a GA with multiple-parent crossover, but with a potentially large number of parents.

Cat Swarm Optimization, CSO, Chu et al., 2006, ~300 citations

---

PSO-style algorithm. At each iteration, each particle either carries out a local search, or moves in the direction of the population best. When carrying out a local search, a specified number of points are sampled in the vicinity of the current position and the best one is kept.

Charged System Search, CSS, Kaveh and Talatahari, 2010, ∼600 citations

PSO-style algorithm. All particles are attracted towards one another. The degree of attraction is calculated using an inverse-square law weighted by fitness. Particles have velocities, but there is no notion of personal best. Adaptive parameter changes allow the degree of attraction to vary over time, and elitism prevents the loss of the global best solutions.

Chemical Reaction Optimization, CRO, Lam and Li, 2010, >300 citations

Population members carry out either local search or a more disruptive global search using disruptive operators such as EA-style mutation and crossover. The balance between local and global search, and the likelihood of accepting lower fitness solutions, are both based on the history of the population member: if no improvement has been made for a while, global search replaces local search; if lower fitness solutions were previously accepted, then they are less likely to be accepted in the future. Similarities to memetic algorithms and simulated annealing.

Cuckoo Optimization Algorithm, COA, Rajabioun, 2011, ∼500 citations

PSO-style algorithm. At each iteration, new search points are sampled within a radius of each existing search point, and only the best points are kept. The resulting population is then clustered using k-means clustering, and the cluster with the highest mean fitness is identified. Search points in the other clusters are then moved towards the fittest cluster.

Cuckoo Search, CS, Yang and Deb, 2009, ∼3000 citations

Uses a small population of solutions. At each iteration, a poor solution is replaced either randomly or by using a 'Lévy flight' applied to another, randomly selected, solution. Lévy flights are a kind of random walk with step sizes generated from a heavy-tailed probability distribution.

Firefly Algorithm, FA, Yang, 2009, >2000 citations

PSO-style algorithm. All particles are attracted towards one another. The degree of attraction is calculated using an inverse-square law.

Firework Algorithm, FWA, Tan and Zhu, 2010, >300 citations

At each iteration, solutions are sampled in a neighbourhood around each of the fittest solutions in the population. New solutions replace existing solutions only if they are fitter. Neighbourhoods are sampled using a Gaussian distribution centred around the current point. The width of the distribution is inversely proportional to the population best, causing increased intensification as search progresses.

Flower Pollination Algorithm, FPA, Yang, 2012, >500 citations

PSO-style algorithm. At each iteration, particles move either towards the global best or a randomly selected particle. In the former case, the step size is determined by sampling a Lévy distribution (see CS).

Fruit Fly Optimization Algorithm, FOA, Pan, 2012, >600 citations

PSO-style algorithm. All particles move towards the current global best. However, how this is achieved is unclear from the description.

Glowworm Swarm Optimization, GwSO, Krishnanand and Ghose, 2005, >600 citations

PSO-style algorithm. Each particle maintains a numerical value that summarises its recent search progress, increasing this when it finds a relatively fit search point and decreasing gradually when it makes no progress. At each iteration, each particle moves towards another particle located within a hyper-spherical region centred around its current point; a particle with a high search progress value is more likely to be followed, and the radius of this region shrinks when there are many particles nearby. Citation count includes (Krishnanand and Ghose, 2009).

Gravitational Search Algorithm, GSA, Rashedi et al., 2009, >2500 citations

PSO-style algorithm. All particles are attracted towards one another. The degree of attraction is calculated using an inverse-square law weighted by fitness. Particles have velocities, but there is no notion of personal best.

Grey Wolf Optimizer, GWO, Mirjalili et al., 2014, >1000 citations

PSO-style algorithm. At each iteration, each solution moves around the edges of a hypercube centred around a target search point. The target point is selected from a region bounded by the three current best search points. Hypercubes become gradually smaller at each iteration in order to intensify search, and there is a random component in the update equation to inject diversity.

Group Search Optimizer, GSO, He et al., 2009, >500 citations

PSO-style algorithm. Particles with highest fitness explore nearby search points, using a mathematical model of animal vision to delimit the region they explore at a particular time. The majority of the other particles move progressively towards the particles with highest fitness. The remaining particles generate diversity by carrying out random walks.

Harmony Search, HS, Geem et al., 2001, ∼4000 citations

At each iteration, a single new solution is created from a randomly selected existing solution. For each of its decision variables, a new value is chosen either at random or by copying and slightly mutating the value from another randomly-selected solution. If the new solution is fitter than the least fit solution in the population, it replaces it. Note that this algorithm has been proven equivalent to a certain form of ES (Weyland, 2015).

Imperialist Competitive Algorithm, ICA, Atashpaz-Gargari and Lucas, 2007, ∼1500 citations

A population is randomly initialised and the fittest n solutions are selected. For each of these solutions, a sub-population is created with size proportional to fitness and is filled randomly using the remaining population members. At each iteration, the solutions in the sub-population are moved, using a PSO-like operation, towards their fittest member, with some noise added. Then, each sub-population is given a value based mainly upon the fitness of its fittest member, and solutions in sub-populations with low values are re-allocated to sub-populations with high values. The algorithm terminates when there is a single non-empty sub-population.

Intelligent Water Drops, IWD, Shah-Hosseini, 2009, >300 citations

ACO-style algorithm for combinatorial optimisation. Agents sequentially follow paths around the solution space, and use real-valued markers to indicate components which were involved in the construction of high fitness solutions. Unlike ACO, the amount that an agent changes a marker is also affected by the values of markers it has passed on its path so far. This means that agents which are following well-marked paths tend to have a greater effect on markers.

Invasive Weed Optimization, IWO, Mehrabian and Lucas, 2006, >750 citations

At each iteration, population members produce offspring in proportion to their fitness using a mutation operator that uses a normal distribution whose width decreases non-linearly over the course of a run. Once the population size reaches an upper bound, low-fitness species (groups of parents and their children) are removed from the population. Similar to GA and ES with distribution-based mutations.

Krill Herd, KH, Gandomi and Alavi, 2012, >600 citations

PSO-style algorithm. Particles are attracted towards the population best, their personal best, and the population "centre of mass" (a fitness-weighted average). They are also either attracted to or repelled by other search points within a given radius, based upon their fitness. Search points also have a component of random motion. The weighting of components is time-dependent, with less random motion and more movement towards the global best as time proceeds. On top of this, fitness-proportionate crossover and mutation operators are applied.

Marriage in Honey Bees Optimization, MBO, Abbass, 2001, ∼400 citations

At each iteration, a number of random walks are carried out, starting from the locations of the fittest n solutions in the population. New solutions are created using a crossover operator that recombines the existing (start of walk) solution with solutions encountered during the walk. The likelihood of this occurring at each step of the walk is based on the fitnesses of the two parents, and also reduces over the course of the walk. Step size progressively decreases during the walk. Local search is used to improve solutions at each iteration of the algorithm; the operator used for this is chosen probabilistically based on its past success rate.

Moth-Flame Optimization, MFO, Mirjalili, 2015a, ∼250 citations

PSO-style algorithm. Particles move in spiral paths towards a target position. The target positions are the personal bests of other particles. Initially, all personal bests are used as targets, with the particle that has the highest current fitness moving towards the fittest personal best, and the particle with the least current fitness moving towards the least fit. Over time, fewer targets are followed.

Shuffled Frog Leaping Algorithm, SFLA, Eusuff and Lansey, 2003, >1000 citations

At each iteration, the population is divided into sub-populations, each with a broad fitness spread. Each sub-population is then repeatedly sub-sampled using a fitness-proportional operator, and the least fit individual in each sub-sample is mutated towards the fittest individual in the sub-sample (or alternatively the population), using a PSO-like operator. In each case, if this does not lead to an improvement in fitness, the individual is reinitialised. After each sub-population has been processed, the sub-populations are merged, and the procedure is repeated.

Society and Civilisation Algorithm, SCA, Ray and Liew, 2003, >300 citations

At each iteration, the population is clustered. In each cluster, the fittest solutions are selected. The remaining solutions in the cluster are then moved towards the selected solutions. A similar procedure is then carried out for the selected solutions from all clusters, with the less fit of these moved towards the fittest. The algorithm uses a rank-based procedure to identify fit solutions, taking into account constraint satisfaction in addition to objective values.

Teacher-Learning Based Optimization, TLBO, Rao et al., 2011, >1000 citations

PSO-style algorithm. At each iteration, the mean position of all particles is calculated and subtracted from the population best. Moves are then carried out by adding a fraction of the resulting vector to each particle (this move is similar to differential evolution). Only improving moves are accepted. Each particle is then compared to another randomly selected target particle; if the target is fitter, it moves towards it; otherwise, it moves away. Again, only improving moves are accepted.

Water Cycle Algorithm, WCA, Shah-Hosseini, 2009, ∼250 citations

PSO-style algorithm. At each iteration, n particles with highest fitness (but not the population best) are moved closer to the population best by a random amount. The remaining less fit particles are each moved closer to one of these n particles by a random amount, with proportionally more of them moving towards the fitter particles among the n. A mutation operator is also applied to maintain diversity.

Whale Optimization Algorithm, WOA, Mirjalili and Lewis, 2016, ∼250 citations

PSO-style algorithm. Each particle moves in a hypercube around a target search point and iteratively moves towards this target either by shrinking the hypercube or through a spiral motion. Target choice is affected by a time-dependent parameter; initially this causes random members of the population to be followed; later all search points follow the population best.

3 Comments on Novelty

Recent nature-inspired metaheuristics have sometimes been criticised for a lack of novelty. Before discussing this in more detail, it is first useful to consider the meaning of the term *metaheuristic*. Many authors who develop nature-inspired algorithms use this term as a synonym for "optimisation algorithm", but this is not the original meaning of the term, which is more akin to a generative model that can be used to guide the development of a particular algorithm. Sörensen et al. (2018) address this disparity by distinguishing metaheuristic algorithms (i.e. particular implementations of a metaheuristic idea) from metaheuristic frameworks (i.e. the more general models from which these algorithms are

derived). This distinction is important when talking about novelty, because whilst there is considerable scope for designing a novel metaheuristic algorithm, there is much less scope for developing a novel metaheuristic framework. For instance, you can create a novel metaheuristic algorithm by modifying the mutation operator used by a GA, or by hybridising a GA with an operator from PSO, but in both cases there is no novel metaheuristic framework being created. It is worth noting that hybridisation, in particular, introduces combinatorial scope for generating algorithms that are technically novel, yet which introduce no novel algorithmic features.

Whilst metaheuristic frameworks are a useful concept for narrowing the definition of novelty, it can also be useful to talk about recurring ideas that appear within multiple frameworks. For instance, EAs and PSO are probably good candidates for being called metaheuristic frameworks, but there are clearly common concepts that occur within both of these; for example, the way in which both techniques have mechanisms for exploring search points that are intermediate to existing ones. In a previous paper (Lones, 2014), I attempted to identify and describe some of these more general metaheuristic approaches; an abridged listing of these is reproduced in Table 1.

Technically, almost all the nature-inspired algorithms described in the previous section meet the definition of a novel metaheuristic algorithm, since they differ from standard metaheuristic algorithms such as ESs, GAs and standard PSO. However, it is difficult to argue that any of them are novel metaheuristic frameworks, since most of them clearly borrow (or perhaps re-discover) concepts that are also central to conventional metaheuristic frameworks. Referring to the metaheuristic concepts listed in Table 1, all of the algorithms described in the previous section implement a combination of *hill climbing*, *adaptive memory programming* and *population-based search*, and this is also true of EAs and PSO. The majority also implement some form of *intermediate search*, most commonly using either a PSO-like operator that picks a point geometrically between two existing points or an EA-like crossover operator that recombines solution components. Those which use PSO-like operators also carry out *directional search* in a similar manner to PSO. Many of the algorithms use *restarts* (ABC, BFO, BeA, CS, SFLA), which are also commonly used in local search algorithms. Many also have strategies for *accepting negative moves*: some of these resemble simulated annealing (BA, CRO); however, the most common approach involves random walks (ALO, BFO, CS, GSO, KH, MBO), which might be considered a degenerate form of threshold acceptance, but is otherwise a relatively novel idea. Several algorithms use search trajectories that follow a spiral-like path around local optima (GWO, MFO, WOA), and this could be considered a form of *variable neighbourhood search*. IWD carries out *search space mapping* in an ACO-like manner.

In terms of resemblance to existing metaheuristic frameworks, PSO-style algorithms (ABC, BeA, BA, COA, CSO, CSS, FA, FOA, FPA, GSA, GSO, GWO, GwSO, KH, MFO, TLBO, WCA, WOA) are by far the most common in the list. A number of algorithms not explicitly tagged as PSO-style also explore ideas that would be familiar to the PSO community (ALO, BFO, BB-BC, FWA, ICA, SFLA, SCA). There is only one ACO-style algorithm mentioned, though other less-cited examples can be found in the literature (e.g. monkey search (Mucherino and Seref, 2007)). A few algorithms

Concept	Description	Examples
Hill Climbing	Follow a sequence of local improvements to reach a locally optimal solution.	Steepest ascent, stochastic hill climbing
Accepting Negative Moves	Allow moves to worse solutions.	Threshold accepting, simulated annealing
Restarts	Restart the search process in a different region once it has converged at a local optimum.	Random-restart hill climbing, iterated local search
Adaptive Memory Programming	Use memory of past search experience to guide future search.	Tabu search, EAs, PSO
Population-Based Search	Multiple cooperating search processes that run in parallel.	EAs, PSO, scatter search
Intermediate Search	Explore the region between two or more previously visited search points.	Crossover, PSO, path relinking
Directional Search	Identify productive directions within the search space, and carry out moves accordingly.	Gradient ascent, CMA-ES, PSO
Variable Neighbourhood Search	Search different neighbourhoods around the location of a known local optimum.	PSO, variable neighbourhood search
Search Space Mapping	Construct a map to guide search processes that are traversing the search space.	ACO, guided local search, DIRECT

Table 1: A list of recurring metaheuristic concepts. Adapted from Lones (2014).

might be considered variants of ES (BeA, HS, IWO), and a number of algorithms are broadly EA-like (BBO, COA, ICA, SFLA, SCA), with a number of these hybridising PSO-like operators (COA, ICA, SFLA, SCA). Some algorithms have notable degrees of self-similarity: for instance, CSS, FA and FSA all use inverse-square laws to calculate the attraction between particles.

4 PSO-Style Algorithms

The majority of the algorithms in the list are either tagged as being PSO-style or use PSO-like operators. Reading through the descriptions in Section 2, it is fairly evident that these algorithms differ in many ways from canonical forms of PSO. However, it should also be borne in mind that PSO has been an active area of research for over 20 years, and during

this time many variants have been developed (for a recent review of these, see (Bonyadi and Michalewicz, 2017)). In effect, development within the PSO and "nature-inspired algorithm" communities has occurred in parallel, with limited interaction between them. Consequently, it is important to gain an understanding of how ideas explored within this group of algorithms intersect with those explored in the PSO literature.

Before attempting to do this, first of all a quick summary of standard PSO: At each iteration, all particles move towards other points. These points are the personal best of the particle and the best point seen by a group of other particles, known as informants. Informants are allocated to each particle at the start of a run, and are often chosen randomly. Elitism can be added by also using the population best as an informant. A particle's next move is computed by adding a velocity vector, which is a weighted sum of the particle's current velocity (which is random in the first iteration) and its current deviation from its informants' best seen point, with a degree of stochasticity added. Taking into account the previous velocity gives the particles momentum, allowing them to overshoot their target points by differing amounts. This is thought to be an important source of diversification in PSO. However, it is also a source of complexity, and makes it difficult to analyse the behaviour of the algorithm.

The majority of PSO-style algorithms listed in Section 2 have similar basic mechanics to PSO, in that particles are moved towards informants using vector operations. A major difference is that the majority of these algorithms (all except KH and MFO) do not use personal bests. This means that particles are not influenced by their own previous locations, and they only move towards the current locations of their informants. The metaheuristic motivation for this is unclear, since it appears to reduce the amount of information available to guide search. Nevertheless, the idea of "social-only" interactions (i.e. ignoring a particle's own search experience) has also been explored in PSO and in both (Kennedy, 1997) and (Pedersen and Chipperfield, 2010) was found to have no significant effect upon the algorithm's performance.

Another major difference from standard PSO is that most of the algorithms have no direct analogue of velocity or momentum; rather, move sizes are determined using simpler rules, including time-dependent move sizes (ABC, ALO, BB-BC, CSS, IWO, MBO), distance-dependent move sizes (CSS, FA, GSA) and region-based sampling (BA, BeA, BB-BC, FWA, GWO, WOA). Time-dependent move sizes have also been explored in variants of PSO (Ratnaweera et al., 2004; Shi and Eberhart, 1999). Region-based sampling, rather than applying vector operations to the current particle location, involves directly sampling from a region of search space that is shaped or bounded by one or more informants. This approach has earlier been used in Bare Bones PSO (Kennedy, 2003), where it was introduced as a means of simplifying the dynamics of PSO and making it more tractable for analysis. Distance-dependent move sizes are notable: usually in PSO, particles move faster towards informants that are further away, meaning that move size increases with distance. In CSS, FA and GSA, on the other hand, particles are less influenced by distant particles, so move size reduces with distance. This causes interactions between particles to become geographically localised, which could be useful for multi-modal landscapes; however, it is unclear whether the resulting behaviour is more effective than other mechanisms introduced to PSO to handle these kind of landscapes,

e.g. multi-swarm approaches (Blackwell and Branke, 2004).

A consequence of using simpler update rules is that the particle dynamics of many of these algorithms are much simpler than in standard PSO. The benefit of this is that it makes the behaviour easier to understand. However, by removing exploratory dynamics like overshooting and oscillation, there is a danger that they will only explore the regions between existing search points and suffer premature convergence as a result. To address this, most include one or more mechanisms to promote diversification. These include hybrididation with local search (CSO, CRO, COA, FWA, IWO, MBO, WCA), random restarts (ABC, BFO, BeA, CS, SFLA), random walks (ALO, BFO, CS, GSO, KH, MBO) and spiral-like movements (GWO, MFO, WOA). The latter, in particular, may lead to particle trajectories that resemble those seen in PSO (and it should be noted that a similar approach is used in spiral optimisation (Tamura and Yasuda, 2011)). Hybridisation with local search is also fairly common in PSO, e.g. (Chen et al., 2005), where restarts have also been used (Kaucic, 2013). Random walks are arguably one of the more interesting mechanisms explored in recent nature-inspired optimisation algorithms, particularly those that build upon biological knowledge in this area, e.g. CS and BFO, and there is no real analogue in the PSO literature.

The manner of choosing informants varies widely amongst the PSO-style algorithms in the list. Some (BA, CSO, FOA) only use the population best, relying on other mechanisms (e.g. restarts) to maintain diversity. Several algorithms (GwSO, WCA, GWO, COA) choose informants in a fitness-informed manner, either probabilistically, by selecting the top n solutions in the population, or in the case of COA, by clustering and identifying the cluster with the highest mean fitness as an informant. These approaches are somewhat related to variants of PSO that use dynamic allocation of informants, e.g. (Du et al., 2015). A number of algorithms have mechanisms that cause particles to be more influenced by nearby informants. This includes those that relate move size to distance (see above). It also includes SCA, which dynamically clusters the population based on distance. Distance-based selection of informants has also been used in PSO, e.g. (Lane et al., 2008). A number of algorithms use all other particles as informants, either directly (CSS, FA, GAO), or indirectly by summarising information about them (KH, BB-BC). Similar ideas have been investigated in fully-informed PSO (Mendes et al., 2004). Some algorithms use time-varying rules for choosing informants, notably those that move from randomly-chosen informants towards the population best over time (MFO, WOA). The idea of dynamically-varying the number of informants over time has also been explored in the PSO literature (Suganthan, 1999).

5 Conclusions

Are recent nature-inspired algorithms *novel*? Yes and no. On the one hand, most (but certainly not all) of the algorithms reviewed in this paper are distinct from existing optimisation algorithms, and given a particular search space, they would likely follow different trajectories to existing algorithms. On the other hand, many of these algorithms use variants of well-established metaheuristics that are also found in existing metaheuristic

frameworks such as PSO, EAs and local search. Furthermore, the analysis of PSO-style algorithms shows that many of their underlying ideas have also been explored by the more mainstream PSO community. However, chronologically, this hasn't always been in one direction. Sometimes the PSO community has explored these ideas earlier, sometimes later, and sometimes in parallel to recent nature-inspired algorithms. Either way, it shows how the fragmentation of the nature-inspired computing community has led to duplicated effort.

Are recent nature-inspired algorithms *competitive*? This is less clear. Most of the cited papers include a performance evaluation. The results are not reported here, because almost all show the algorithm to perform better than the algorithms they were compared against. Even without taking No Free Lunch theorems (Wolpert and Macready, 1997) into account, it is implausible to believe that this is true for all of them. This is not to say that the results are incorrect, but it does reflect the difficulty of designing fair comparative studies (Črepinšek et al., 2016; Fong et al., 2016; García-Martínez et al., 2017; Piotrowski, 2015). We can speculate that all of these algorithms will *sometimes* perform better on *some* problems when compared against other algorithms, since problem landscapes are diverse, and small differences in the topography of a landscape can favour different approaches.

However, given a *specific* problem, it is difficult to know which algorithm will work well. The field of meta-learning (Lemke et al., 2015) has been studying this issue for some time, but progress on understanding how problems can be characterised, categorised and mapped to specific optimisers has so far been limited. This means that performance on one problem currently tells us little about potential performance on another problem, and consequently that practitioners usually have to try out a range of different optimisers to determine which one works well on their problem. In a sense, the recent developments in nature-inspired algorithms have increased the number of optimisers available to try out. This may sometimes be beneficial, but it also makes it harder for a practitioner to identify a suitable optimiser that is well-understood and has community support. Given the vast scope for creating variants and hybrids of existing algorithms, this situation is only likely to get worse.

An alternative, and arguably more promising, direction of travel can be seen in the hyperheuristics (Burke et al., 2013; Epitropakis and Burke, 2018) and broader machine learning communities (Li and Malik, 2017; Wichrowska et al., 2017). Both address the problem of choosing an optimiser as an optimisation problem, using a machine learning algorithm to identify an optimiser that is good at solving a specific task. In the case of hyperheuristics, the optimiser, which is usually an evolutionary algorithm, can be used to construct new optimisation algorithms. This can be done either by specialising an existing algorithm (for example, evolving a new mutation operator for an EA) or by assembling existing metaheuristic components in a novel way. In effect, the latter is an automated version of the many manual attempts to hybridise metaheuristics that can be found in the literature. However, this automated approach is currently limited by a lack of standardised interfaces (Swann and Hammond, 2015), and this arguably is limited by the tendency of the community to think of metaheuristics in terms of algorithms rather than re-usable components.

This focus on algorithms rather than components is a particular issue for the nature-inspired algorithm community, where the objective of domain modelling is almost always the generation of a single algorithm that captures all pertinent behaviours present within the domain of inspiration. As a consequence, any interesting, novel, components extracted from the domain tend to become conflated with other, less interesting, and sometimes arbitrary, components. This makes it hard to understand the relevance and contribution of individual components within the optimisation setting. Arguably a better approach would be to identify any component of the domain that is particularly interesting, and integrate this individually within one or more existing metaheuristic frameworks. Even better would be to make the code available in re-usable form: it could then be used by other algorithm developers, or even used as a new building-block within hyperheuristic frameworks.

References

H. A. Abbass. MBO: Marriage in honey bees optimization—a haplometrosis polygynous swarming approach. In *Proceedings of the 2001 Congress on Evolutionary Computation (CEC 2001)*, volume 1, pages 207–214. IEEE, 2001.

E. Atashpaz-Gargari and C. Lucas. Imperialist competitive algorithm: an algorithm for optimization inspired by imperialistic competition. In *Proceedings of the 2007 Congress on Evolutionary Computation (CEC 2007)*, pages 4661–4667. IEEE, 2007.

T. Blackwell and J. Branke. Multi-swarm optimization in dynamic environments. In *Workshops on Applications of Evolutionary Computation*, pages 489–500. Springer, 2004.

M. R. Bonyadi and Z. Michalewicz. Particle swarm optimization for single objective continuous space problems: a review, 2017.

E. K. Burke, M. Gendreau, M. Hyde, G. Kendall, G. Ochoa, E. Özcan, and R. Qu. Hyper-heuristics: A survey of the state of the art. *Journal of the Operational Research Society*, 64(12):1695–1724, 2013.

J. Chen, Z. Qin, Y. Liu, and J. Lu. Particle swarm optimization with local search. In *International Conference on Neural Networks and Brain (ICNN&B'05)*, volume 1, pages 481–484. IEEE, 2005.

S.-C. Chu, P.-W. Tsai, and J.-S. Pan. Cat swarm optimization. In *Pacific Rim International Conference on Artificial Intelligence*, pages 854–858. Springer, 2006.

M. Črepinšek, S.-H. Liu, and L. Mernik. A note on teaching–learning-based optimization algorithm. *Information Sciences*, 212:79–93, 2012.

M. Črepinšek, S.-H. Liu, L. Mernik, and M. Mernik. Is a comparison of results meaningful from the inexact replications of computational experiments? *Soft Computing*, 20(1): 223–235, 2016.

W.-B. Du, Y. Gao, C. Liu, Z. Zheng, and Z. Wang. Adequate is better: particle swarm optimization with limited-information. *Applied Mathematics and Computation*, 268: 832–838, 2015.

M. G. Epitropakis and E. K. Burke. Hyper-heuristics. *Handbook of Heuristics*, pages 1–57, 2018.

O. K. Erol and I. Eksin. A new optimization method: big bang–big crunch. *Advances in Engineering Software*, 37(2):106–111, 2006.

M. M. Eusuff and K. E. Lansey. Optimization of water distribution network design using the shuffled frog leaping algorithm. *Journal of Water Resources Planning and Management*, 129(3):210–225, 2003.

I. Fister Jr, X.-S. Yang, I. Fister, J. Brest, and D. Fister. A brief review of nature-inspired algorithms for optimization. *Elektrotehniški vestnik*, 80(3), 2013.

S. Fong, X. Wang, Q. Xu, R. Wong, J. Fiaidhi, and S. Mohammed. Recent advances in metaheuristic algorithms: Does the Makara dragon exist? *The Journal of Supercomputing*, 72(10):3764–3786, 2016.

A. H. Gandomi and A. H. Alavi. Krill herd: a new bio-inspired optimization algorithm. *Communications in Nonlinear Science and Numerical Simulation*, 17(12):4831–4845, 2012.

C. García-Martínez, P. D. Gutiérrez, D. Molina, M. Lozano, and F. Herrera. Since CEC 2005 competition on real-parameter optimisation: a decade of research, progress and comparative analysis's weakness. *Soft Computing*, 21(19):5573–5583, 2017.

Z. W. Geem, J. H. Kim, and G. V. Loganathan. A new heuristic optimization algorithm: harmony search. *simulation*, 76(2):60–68, 2001.

S. He, Q. H. Wu, and J. Saunders. Group search optimizer: an optimization algorithm inspired by animal searching behavior. *IEEE Transactions on Evolutionary Computation*, 13(5):973–990, 2009.

D. Karaboga. An idea based on honey bee swarm for numerical optimization. Technical report, Technical report-tr06, Erciyes University, engineering faculty, computer engineering department, 2005.

M. Kaucic. A multi-start opposition-based particle swarm optimization algorithm with adaptive velocity for bound constrained global optimization. *Journal of Global Optimization*, 55(1):165–188, 2013.

A. Kaveh and S. Talatahari. A novel heuristic optimization method: charged system search. *Acta Mechanica*, 213(3-4):267–289, 2010.

J. Kennedy. The particle swarm: social adaptation of knowledge. In *IEEE International Conference on Evolutionary Computation*, pages 303–308. IEEE, 1997.

J. Kennedy. Bare bones particle swarms. In *Proceedings of the 2003 IEEE Swarm Intelligence Symposium (SIS'03)*, pages 80–87. IEEE, 2003.

K. Krishnanand and D. Ghose. Detection of multiple source locations using a glowworm metaphor with applications to collective robotics. In *Proceedings 2005 IEEE Swarm Intelligence Symposium (SIS 2005)*, pages 84–91. IEEE, 2005.

K. Krishnanand and D. Ghose. Glowworm swarm optimization for simultaneous capture of multiple local optima of multimodal functions. *Swarm intelligence*, 3(2):87–124, 2009.

A. Y. Lam and V. O. Li. Chemical-reaction-inspired metaheuristic for optimization. *IEEE Transactions on Evolutionary Computation*, 14(3):381–399, 2010.

J. Lane, A. Engelbrecht, and J. Gain. Particle swarm optimization with spatially meaningful neighbours. In *Proceedings 2008 IEEE Swarm Intelligence Symposium (SIS 2008)*, pages 1–8. IEEE, 2008.

C. Lemke, M. Budka, and B. Gabrys. Metalearning: a survey of trends and technologies. *Artificial intelligence review*, 44(1):117–130, 2015.

K. Li and K. Malik. Learning to optimize. In *5th International Conference on Learning Representations*, 2017.

M. A. Lones. Metaheuristics in nature-inspired algorithms. In *Proceedings of the Companion Publication of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 1419–1422. ACM, 2014.

A. R. Mehrabian and C. Lucas. A novel numerical optimization algorithm inspired from weed colonization. *Ecological informatics*, 1(4):355–366, 2006.

R. Mendes, J. Kennedy, and J. Neves. The fully informed particle swarm: simpler, maybe better. *IEEE Transactions on Evolutionary Computation*, 8(3):204–210, 2004.

S. Mirjalili. Moth-flame optimization algorithm: A novel nature-inspired heuristic paradigm. *Knowledge-Based Systems*, 89:228–249, 2015a.

S. Mirjalili. The ant lion optimizer. *Advances in Engineering Software*, 83:80–98, 2015b.

S. Mirjalili and A. Lewis. The whale optimization algorithm. *Advances in Engineering Software*, 95:51–67, 2016.

S. Mirjalili, S. M. Mirjalili, and A. Lewis. Grey wolf optimizer. *Advances in Engineering Software*, 69:46–61, 2014.

A. Mucherino and O. Seref. Monkey search: a novel metaheuristic search for global optimization. In *AIP Conference Proceedings*, volume 953, pages 162–173, 2007.

W.-T. Pan. A new fruit fly optimization algorithm: taking the financial distress model as an example. *Knowledge-Based Systems*, 26:69–74, 2012.

K. M. Passino. Biomimicry of bacterial foraging for distributed optimization and control. *IEEE Control Systems*, 22(3):52–67, 2002.

M. E. H. Pedersen and A. J. Chipperfield. Simplifying particle swarm optimization. *Applied Soft Computing*, 10(2):618–628, 2010.

D. T. Pham, A. Ghanbarzadeh, E. Koç, S. Otri, S. Rahim, and M. Zaidi. The bees algorithm—a novel tool for complex optimisation problems. In *Intelligent Production Machines and Systems*, pages 454–459. Elsevier, 2006.

A. P. Piotrowski. Regarding the rankings of optimization heuristics based on artificially-constructed benchmark functions. *Information Sciences*, 297:191–201, 2015.

R. Rajabioun. Cuckoo optimization algorithm. *Applied soft computing*, 11(8):5508–5518, 2011.

R. V. Rao, V. J. Savsani, and D. Vakharia. Teaching–learning-based optimization: a novel method for constrained mechanical design optimization problems. *Computer-Aided Design*, 43(3):303–315, 2011.

E. Rashedi, H. Nezamabadi-Pour, and S. Saryazdi. GSA: a gravitational search algorithm. *Information sciences*, 179(13):2232–2248, 2009.

A. Ratnaweera, S. K. Halgamuge, and H. C. Watson. Self-organizing hierarchical particle swarm optimizer with time-varying acceleration coefficients. *IEEE Transactions on Evolutionary Computation*, 8(3):240–255, 2004.

T. Ray and K. M. Liew. Society and civilization: An optimization algorithm based on the simulation of social behavior. *IEEE Transactions on Evolutionary Computation*, 7(4):386–396, 2003.

H. Shah-Hosseini. The intelligent water drops algorithm: a nature-inspired swarm-based optimization algorithm. *International Journal of Bio-Inspired Computation*, 1(1-2): 71–79, 2009.

Y. Shi and R. C. Eberhart. Empirical study of particle swarm optimization. In *Proceedings of the 1999 Congress on Evolutionary Computation (CEC 99)*, volume 3, pages 1945–1950. IEEE, 1999.

D. Simon. Biogeography-based optimization. *IEEE transactions on evolutionary computation*, 12(6):702–713, 2008.

K. Sörensen. Metaheuristics—the metaphor exposed. *International Transactions in Operational Research*, 22(1):3–18, 2015.

K. Sörensen, M. Sevaux, and F. Glover. A history of metaheuristics. *Handbook of heuristics*, pages 1–18, 2018.

P. N. Suganthan. Particle swarm optimiser with neighbourhood operator. In *Proceedings of the 1999 Congress on Evolutionary Computation (CEC 99)*, volume 3, pages 1958–1962. IEEE, 1999.

J. Swann and K. Hammond. Towards 'metaheuristics in the large'. In *Proceedings of 11th Metaheuristics International Conference (MIC 2015)*, 2015.

K. Tamura and K. Yasuda. Primary study of spiral dynamics inspired optimization. *IEEJ Transactions on Electrical and Electronic Engineering*, 6(S1), 2011.

Y. Tan and Y. Zhu. Fireworks algorithm for optimization. In *International Conference in Swarm Intelligence*, pages 355–364. Springer, 2010.

D. Weyland. A critical analysis of the harmony search algorithm—how not to solve sudoku. *Operations Research Perspectives*, 2:97–105, 2015.

O. Wichrowska, N. Maheswaranathan, M. W. Hoffman, M. D. Sergio Gómez Colmenarejo, N. Freitas, and J. Sohl-Dickstein. Learned optimizers that scale and generalize. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, 2017.

D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, 1(1):67–82, 1997.

B. Xing and W.-J. Gao. *Innovative computational intelligence: a rough guide to 134 clever algorithms.* Springer, 2016.

X.-S. Yang. Firefly algorithms for multimodal optimization. In *International symposium on stochastic algorithms*, pages 169–178. Springer, 2009.

X.-S. Yang. *Nature-inspired metaheuristic algorithms.* Luniver press, 2010.

X.-S. Yang. Flower pollination algorithm for global optimization. In *International conference on unconventional computing and natural computation*, pages 240–249. Springer, 2012.

X.-S. Yang and S. Deb. Cuckoo search via lévy flights. In *World Congress on Nature & Biologically Inspired Computing (NaBIC 2009)*, pages 210–214. IEEE, 2009.

X.-S. Yang and A. H. Gandomi. Bat algorithm: a novel approach for global engineering optimization. *Engineering Computations*, 29(5):464–483, 2012.