



HAL
open science

Evolving embodied intelligence from materials to machines

David Howard, Agoston E Eiben, Danielle Frances Kennedy, Jean-Baptiste Mouret, Philip Valencia, Dave Winkler

► **To cite this version:**

David Howard, Agoston E Eiben, Danielle Frances Kennedy, Jean-Baptiste Mouret, Philip Valencia, et al.. Evolving embodied intelligence from materials to machines. *Nature Machine Intelligence*, 2019, 1 (1), pp.12-19. 10.1038/s42256-018-0009-9 . hal-01986599

HAL Id: hal-01986599

<https://inria.hal.science/hal-01986599>

Submitted on 18 Jan 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Evolving embodied intelligence from materials to machines

David Howard^{1*}, Agoston E. Eiben², Danielle Frances Kennedy¹, Jean-Baptiste Mouret³, Philip Valencia¹, and Dave Winkler^{1,4,5}

Natural lifeforms specialise to their environmental niches across many levels; from low-level features such as DNA and proteins, through to higher-level artefacts including eyes, limbs, and overarching body plans. We propose Multi-Level Evolution (MLE), a bottom-up automatic process that designs robots across multiple levels and niches them to tasks and environmental conditions. MLE concurrently explores constituent molecular and material 'building blocks', as well as their possible assemblies into specialised morphological and sensorimotor configurations. MLE provides a route to fully harness a recent explosion in available candidate materials and ongoing advances in rapid manufacturing processes. We outline a feasible MLE architecture that realises this vision, highlight the main roadblocks and how they may be overcome, and show robotic applications to which MLE is particularly suited. By forming a research agenda to stimulate discussion between researchers in related fields, we hope to inspire the pursuit of multi-level robotic design all the way from material to machine.

Perspective article — Howard, David, et al. "Evolving embodied intelligence from materials to machines." *Nature Machine Intelligence* 1.1 (2019): 12.

Robots are on the rise, and seen with increasing ubiquity in what are known as 'structured' environments. Pick and place machines are a good example; their interactions with the environment are predictable and so easily controllable. As a counterpoint, robots consistently struggle in complex, unpredictable 'unstructured' environments^{1,2}. Cataloguing biodiversity in remote areas, searching destroyed buildings for survivors following an earthquake, and exploring labyrinthine cave systems are good examples.

The challenge of embodied intelligence

Natural life thrives in unstructured environments through a specific brand of intelligence known as *embodied cognition*³. Intelligent behaviour emerges from tight coupling between an agent's body, brain, and environment, not solely from the brain. In the taxonomy of philosophy, it opposes the 'I think therefore I am' of Descartes, and a plethora of research to date has shown that the form and function of an agent's physical presence plays an important role in learning, development, and the generation of suitable in-environment behaviour⁴.

Life's ability to produce useful embodiments comes from a free-form evolutionary process where variance occurs across multiple *levels*⁵. Generally speaking, mutations in low-level DNA lead to changes in protein expression, facilitating an emergent process defining the structure and composition of higher-level features including eyes, hands, and limbs, and their placement in body plans. Making robot design similarly free-form and level-based might herald a new wave of capable embodiments to finally tackle challenging unstructured environments.

Embodied cognition and its artificial analogue, embodied Artificial Intelligence (AI)^{6,4}, have long known that complex environments can only be tackled by sufficiently capable combinations of body and brain. In robots, body design has lagged behind due to inherent manufacturing complexities (through cell division, na-

ture gets this ability 'for free'). As such, including a variety of materials into robot design has been a long-standing 'holy grail' for embodied AI⁴ and robotics in general. We present a straightforward, reasonably scaleable path towards incorporating material search and selection into robot design: a more free-form and specialising process than any currently available method.

We call this algorithmic framework Multi-Level Evolution, or MLE. Here we consider a three-level architecture, although as many levels as necessary may be instantiated depending on problem demands. Three levels is a natural split, based on longstanding delineations across the large, active and well-established fields of materials science, robotics, and component design.

At the lowest level, materials are discovered. Components are created by selecting one or more materials into a geometry. Finally, robots are created by integrating components into template 'body plans', and evaluating them on a task in an environment. New candidate materials and components are discovered during the process, continually increasing the range of possible robot designs.

A key takeaway is that MLE directly contrasts conventional engineering approaches which, because of time and cost constraints, often search for versatile generalist robots that do a bit of everything at a reasonable level of performance. MLE is a universal designer (across a wider design space, and for any task-environment niche) that generates specialists by harnessing *diversity* across all levels, and the *emergence* of useful artefacts and artefact combinations.

In this Perspective, we outline how new types of artificial evolution, which are ready to exploit the same wave of ubiquitous computing resources that powered the rise of deep learning, and have been shown to achieve a corresponding leap in performance⁷, can harness the recent explosion of available materials and manufacturing techniques (see boxouts) to create powerfully-embodied robots. We review the main roadblocks to the realisation of this vision. We define the key features of MLE architectures, and, using examples of cutting-edge evolutionary algorithms, propose a simple implementation. We highlight use cases to which MLE is particularly suited. Finally, we sketch out a path towards increasingly capable MLE implementations, and discuss the implications of MLE for the field of robotics.

¹ The Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia

² Vrije Universiteit Amsterdam, Netherlands

³ Inria, CNRS, Université de Lorraine, France

⁴ Latrobe Institute for Molecular Science, La Trobe University, Australia

⁵ Monash Institute of Pharmaceutical Sciences, Monash University Australia

*e-mail: david.howard@csiro.au

Enabling technologies: materials

According to the laws of chemistry, the number of materials available for search is $\approx 10^{100}$ (see ref⁸) (for comparison, there are 'only' 10^{82} atoms in the universe. This provides an almost infinite toolbox for designing bespoke, functional materials for robots, including new sensing, actuation, and power materials⁹, that we are increasingly able to design, characterise, and synthesise. Accelerating the development cycle of bespoke materials is key as the space of possibilities is so vast (using e.g., robotic materials synthesis¹⁰, or combinatorial materials libraries¹¹).

Semi- or fully-autonomous 'closed loop' systems use robots to efficiently perform experiments with reduced reliance on humans, and naturally couples with techniques that automatically plan optimal sets of experiments^{12,13} to reach desired material properties¹⁴. Robots can perform multiple simultaneous experiments, vastly reducing time and human effort whilst increasing providence of experimental data for computational, AI and machine learning methods, which are now mature enough to reliably predict properties of new materials and allow a vast set of previously-physical experiments to be conducted (cheaper and faster) virtually⁸. High-throughput computational techniques exploit massive data-sets and advanced modelling to the same effect^{15,16}.

Combined, these advances provide unprecedented opportunities to design and manufacture 'smarter', more specialised robotic materials that integrate sensing, actuation, and other properties to create powerful embodiment options¹⁷.

Enabling technologies: advanced manufacturing

MLE is poised to exploit additive¹⁸ and subtractive manufacturing of free-form structures and complex geometries¹⁹, printing intricate multi-part components from multiple materials in-situ.

Behavioural diversity can be embedded during the manufacturing process through functional gradation^{20,21} to vary material properties (e.g., stiffness, elasticity). Voxel blending gives fine-resolution, gradual tuning of build properties through continuous mixing of multiple materials. Production devices tuned for ever-expanding range of feed-stock increases the diversity observed in recent composite and multi-material robots, e.g.²².

MLE is iterative, so streamlining construction and reducing reliance on humans is a priority. Sensors, actuators, and power systems are readily printable in various configurations²³, and continue to close the performance gap with their traditional counterparts whilst being increasingly integrated into multi-function materials during construction¹⁷. MLE is poised to benefit from (semi-)autonomous robot construction, including prototype generate-and-test systems^{24,25}, culminating with whole robots constructed without human intervention²⁶.

Inspiration and characteristics of Multi-Level architectures

MLE is a natural extension of Evolutionary Robotics (ER); a field that harnesses iterative, population-based algorithms to generate robot bodies, brains, or both²⁷. A typical ER experiment defines a representation; how the genotype (string of numbers) maps to a phenotype (physical robot). To capture sufficient complexity, these representations are typically indirect - simple genotypes define more complex robot phenotypes, potentially incorporating naturally-observed features including gene reuse (e.g., to encode two identical eyes), radial and bilateral symmetry (seen across nature in body plans), and scaling factors (across the five fingers of a hand)²⁸. It randomly initialises a population using the representation, and tests their task-environment performance against a user defined 'fitness function'. Analogies of genetic mutation and recombination induce variance in the genotypes to create a new generation, with a preference to select high-fitness parents. This process iterates until some acceptable level of performance is met.

ER provides environmental adaptation²⁹, and explores a wider design space than other approaches, locating unconventional 'short-cut' designs which may be otherwise missed³⁰. Owing to a dearth of versatile, affordable manufacturing processes, ER traditionally focused on controller generation for fixed morphologies³¹. Signalled by the first 3D printed evolved robot in 2000³², we now find ourselves in the era of the 'Evolution of Things'^{33,34},

where complex physical artefacts are evolved and physically instantiated³⁵.

From the MLE perspective, 'classic' ER is the top-level level that finds environment and task specific controllers and body plans - arrangements of structure, sensing, and actuation that together comprise a robot. Materials are not typically considered as part of the robot's 'genome' (although idealised materials properties appear sporadically in simulation³⁶). We posit that the missing link to unlocking richer embodiments is to discover, model, and select real (and newly-discovered) materials, and make them available in a holistic design process. MLE architectures are characterised by:

1. Three vertically-stacked levels (robot, component, material). Robots are arrangements of components, where a component is a combination of a geometry and one or more materials that occupy sections of the geometry.
2. At least one search process per level, which is responsible for finding new artefacts within a given level. In the component level, we could run search processes for actuators, sensors, and structural elements.
3. Hybridisation, a novel concept that enforces that real and virtual genomes are identical for either physical or virtual instantiation. This means we can easily 'cross-breed' between simulated and real artefacts.

We suggest evolutionary algorithms as bias-free and domain-agnostic³⁷ default algorithms, with a track record of success in discovering molecules and materials⁸, components and structures³⁸, and robots³⁹, whilst being relatively efficient across all of these levels⁴⁰. As each level is independent, we can use domain specific algorithms/representations as requires. For example, the materials level may benefit from capturing the underlying phenomena relating a materials structure to its behaviour⁴¹.

With the 'grand vision' sketched out, let us now consider how emerging technologies can build simple prototype MLE architectures in the near future.

A conceptual MLE architecture

Our conceptual prototype uses evolutionary Illumination algorithms⁴² (also called Quality-Diversity⁴³) to produce diverse libraries of high-performing potential solutions across three levels (Fig. 1(a)). Libraries are n-dimensional grids of possible combinations of some physical properties, discretised into 'bins'⁴⁴. For example, all actuators transmit force and consume energy, all materials possess weight, rigidity, elasticity, and compliance that can be exploited to generate robots that are adapted to specific environments. By measuring these properties we can assign to the appropriate bin.

For clarity, Fig.1 visualises two properties per process, resulting in 2D grids. Practically, there will be many more properties (Table 1). It is critical that MLE generates a diverse set of solutions, rather than a single optimal solution as in traditional optimization and classic ER. Libraries allow each level to be explored independently and provides diversity to upper levels. Each level may be subdivided; e.g., components can be subdivided into sensors, actuators, and body structure. Each search process can also have its own solution representation, feature dimensions, and search operators.

To begin an experiment, we bootstrap the lowest level with known materials, either from the literature, or from previous MLE experiments. Each material is placed in a bin based on its physical properties. The components level then defines geometries and selects an appropriate number of materials into those geometries to create body segments, sensors, and actuators, filling some component bins. At the highest level, we search for controllers and body plans; templates that define arrangements of

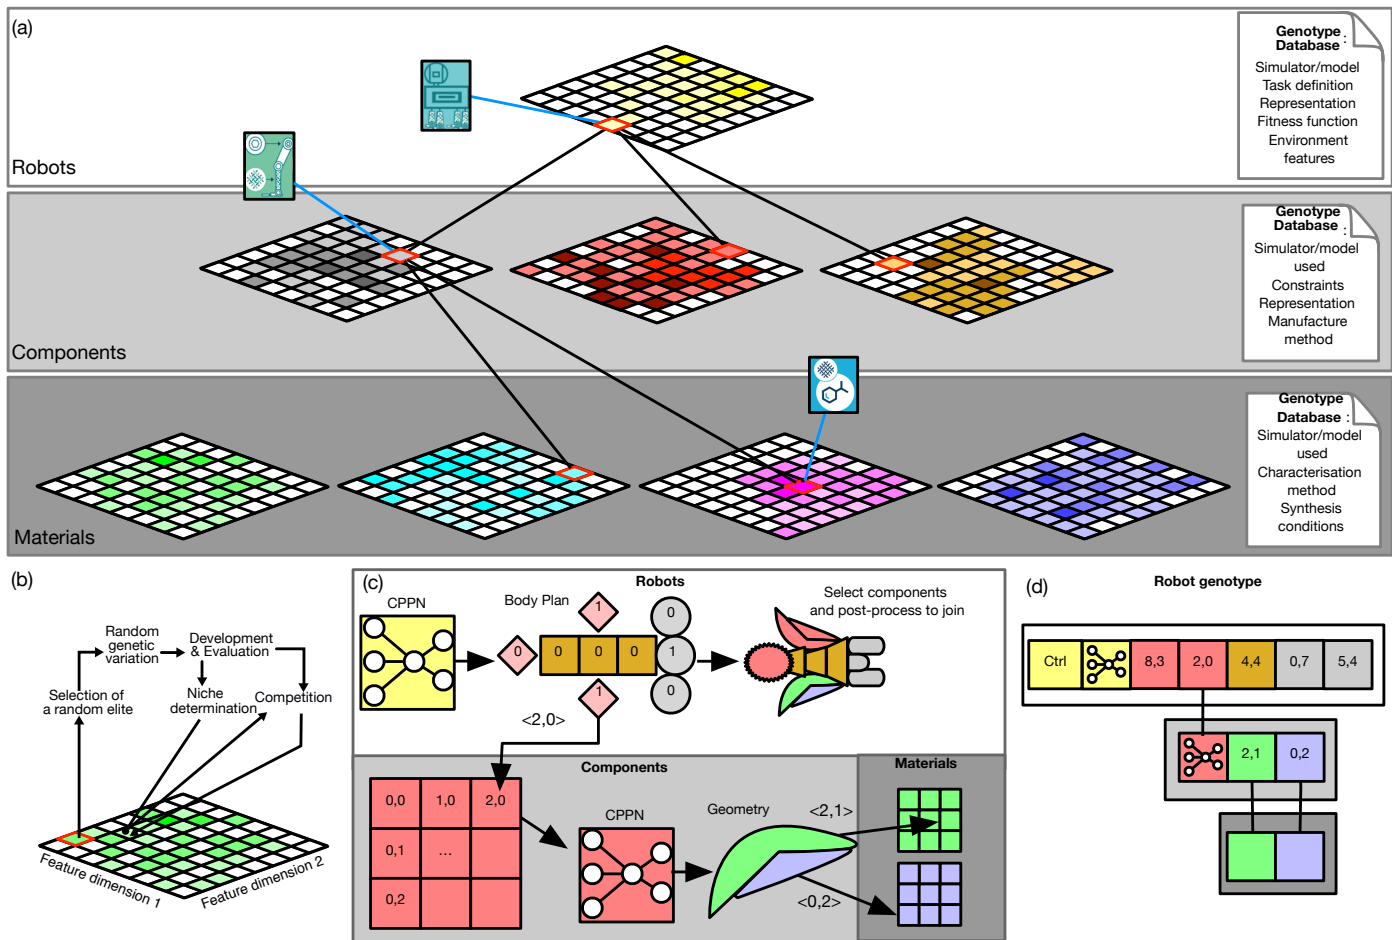


Fig. 1 | (a) A sample MLE architecture incorporating a single robot search process, 3 component search processes (e.g., sensors, actuators, body segments), and 4 materials search processes (polymers, etc.). For clarity only two dimensions of each process are shown, discretised into ‘bins’. Colour indicates the highest fitness solution found per bin; darker represents a fitter solution, white squares indicate no current solution. Over time, more bins are filled, and bin fitness is improved, which improves the quality and diversity of options available to upper levels. An asterisk (*) denotes an individual created physically, other individuals are virtual. **(b) Creating a diversity of high-quality solutions.** At every iteration, the Illumination algorithm (1) randomly selects an occupied bin, (2) adds random mutations to the current best solution of the bin, (3) evaluates the quality (fitness) and the features of the newly generated solution, (4) compares the quality of the newly generated solution with the current best of the bin that corresponds to its features and keeps the best solution. These four steps are repeated until all the bins are occupied with satisfying solutions. **(c) Creating a robot.** At the top level, a Compositional Pattern Producing Network (CPPN) defines the body plan. Once generated, an appropriate number of pointers into the components layer are set based on the number of component slots generated, and the corresponding components fused via post-processing to create the final robot. Pointers address a specific member of the library, in this case two integers for a two-dimensional library. In this case components are segregated into actuation (red diamonds), body structure (gold squares) and sensing (grey circles). A further CPPN per-component defines component geometry, and subsequently the number of material pointers required for that component. Coloured grids (red, green, purple) abstractly represent the larger libraries seen in (a,b). **(d) A hierarchical genotype** that uses ‘pointers’ to fully, efficiently define the robot. The robot controller (‘ctrl’) is defined at the top level. Background colours in (a) (c) and (d) signify different levels.

components. Here, geometries and body plan layouts are defined using Compositional Pattern Producing Networks (CPPNs), specialised neural networks evolved to output geometric patterns displaying modularity, regularity and symmetry, see e.g.⁴⁵.

As well as belonging to a bin, each material, component, and robot has an associated fitness. For materials and components, we suggest fitness based on the universally beneficial property of cost, therefore the cheapest example that fulfils certain physical property requirements will be passed to the next level to reduce the manufacturing burden. More specific fitness measures, e.g., efficiency for an actuator, or signal-to-noise for a sensor, will be mediated by the environment, so we lose transferrability for potential gains in performance. Robot fitness is based on its behaviour; how well it completes the task.

The grids at each level progressively fill out as new material, component, and robot designs are discovered (Fig. 1(b)). Illumination search specifically encourages diversity⁴² through pressure to discover new combinations of physical properties, providing larger libraries and thus more opportunities to exploit ma-

terials and components in interesting ways, facilitating emergent behaviour. As lower levels focus on creating a diversity of options, significant opportunities arise for the spontaneous emergence of component-material combinations that facilitate useful behaviour, which will likely result in a good fitness score for the robot, with no constraints on exactly how that behaviour emerges. These behaviours are a holistic combination of the search efforts at every level, and the in-environment performance of the resulting robot. Counter-intuitively, Illumination search is known to discover more ‘optimal’ outcomes than pure optimisation approaches⁴²; hence we expect high-performance artefacts. Cascading improvements may percolate through the levels; when a new material is found it could improve the fitness of the solution in a populated bin (replacing the previous best), or it may expand the number of filled bins in its level, and potentially the number of reachable bins at any level above it.

To instantiate a robot (1(c)), we query the corresponding CPPN and seamlessly integrate the relevant components into the resulting body plan using post-processing. Accompanying the CPPN

Table 1 | Sample search processes, tunable variables, and desired properties that may be found in each level (non-exhaustive).

Level	Search processes	Tunable variables	Desired properties
Material	Polymer, metal, composite, nanomaterial, hydrogel, thin-film, shape-memory, dielectrics.	Level of experiment, structures and physiochemical properties (of reactants, reagents, substrates, solvents, catalysts, additives, and products), pH, particle size, shape and density, surface roughness and chemistry, pressure, temperature, and time, speed of mixing or coating, rate and order of addition; and post-processing conditions such as drying, conditioning or activation.	Flex, deformation, weight, response strength, energy requirements, hysteresis, signal/noise (sensing), repeatability, power density, peak voltage/current, useful temperature range, cost
Component	Sensor, actuator, body structure, energy, multi-function	Geometric properties, material composition, gradation	Load bearing, strain density, impact resistance, weight, deflection, range of motion, force transmission, torque, compliance, flexibility, sensitivity (to light, chemicals, angular movement, etc...), power requirements, footprint, power provision
Robot	Targeted locomotion, gripping, carrying/transporting, food seeking, recharging	Selected components, arrangement of components in body plan, controllers, post-processing algorithm	Embodied behaviour (from control/robot/environment interactions)

are a number of ‘pointers’ to bins in the components level, which selects specific components into the body plan. Each pointer addresses a specific bin in the level below. Similarly, a component consists of a geometry-defining CPPN and pointers to materials. Either the CPPN or the indices of the pointers may be altered during evolutionary search, which changes the shape or composition of the affected artefacts. Materials may be represented and searched in a similar way. Once instantiated, the robot is evaluated based on desired mission performance to ascertain its fitness.

Unlike natural genotypes, which are defined at the DNA level, the genotype of a robot produced by MLE can be thought of as a hierarchy 1(d), where the robot body plan and controller are defined at the top level. Following the pointers from robot to components allows us to fully define the components used, and following each component’s pointers allow us to fully define the materials that comprise each component.

The only necessary inter-level communication is the passing of candidate solutions upwards for use by the next level. For efficiency, and for ease of integration into higher level simulators/models, only phenotypic properties (i.e., of the physical solution created) are passed between levels. The representation of a solution, plus details of experimental procedures, models/simulators used, learning algorithms, and evaluation tests are stored in a database by the relevant layer as required so results are repeatable.

Not all physical properties will be relevant in all situations, and can be safely ignored. Sparse feature selection methods^{46,47}, applied as automatic dimensionality filters, give more weight to the features most relevant for fitness and function within the niche and minimise combinatorial issues. In our example, this may select the physical properties in which we encourage diversity. Similarly, not all options will be required from lower levels. Each level can also design its own library from lower-level libraries according to its own objectives⁴⁸. Combined, these processes reduce the required computational effort, and the extent of physical characterisation required.

Learning and behaviour

The focus of this Perspective is in improving the bodies of embodied robots. However, we need some way of generating useful behaviour from these bodies. In some cases, this can result solely

from the interactions of materials and components in the robot’s body³⁶, or through automatic response to stresses experienced between body and environment⁴⁹. This *morphological computing* offloads the computation of behaviour from a controller onto the robot’s body⁵⁰, reducing the required controller complexity. We expect MLE to greatly benefit from morphological computing, owing to the vast range of physical behavioural responses it can instantiate.

For more complex tasks, learning will be required to overtly direct the body-environment interactions our robots produce⁵¹, creating a controller — the ‘ctrl’ seen in the robot genome in Fig 1(d). Software provides a wealth of options including neural networks, central pattern generators, behaviour trees, and modular architectures²⁷, which can be optimised through reinforcement learning, evolutionary algorithms, and imitation. Post-deployment *online* learning offers the possibility to adapt controllers following hardware failures⁵², or to changing environmental conditions. Ultimately the choice of controller and learning is a design decision, key requirements are that the body is controllable, its material and morphological composition properly exploited, and its behaviour suitable for our task.

Physical and virtual testing provides the best of both worlds

Manufacturing and testing each new material, component, and robot in reality would be prohibitively expensive in terms of time and cost. MLE’s success hinges on effective use of simulation and modelling, and blurring the lines between real and virtual.

MLE introduces the novel concept of *Hybridisation*, such that the representation describing a material (or component, or robot) is identical, regardless of whether it is real or virtual. Bins in each level may be filled through physical experimentation, or the results of simulation or predictive model. Simulated evolution runs concurrently with physical experimentation, and cross-breeding allows physical or virtual materials (or components, or robots) to parent a child that may exist in the real world, in the virtual world or in both. The advantages of hybridisation are significant; physical evolution is accelerated by the virtual component that can run faster to find good robot features with less time and fewer resources, whereas simulated evolution benefits from the influx of ‘genes’ that are tested favourably in the real world.

Physical experimentation provides necessary ‘ground truth’

data, the burden of which may be reduced through smart algorithmic design⁵³. Evaluating robot performance in reality is particularly difficult (Repeatability and physical damage are key problems), but increasingly feasible thanks to custom-designed test arenas^{54,55}, and proof-of-concept ‘generate and test’ facilities²⁵.

We must be able to simulate the performance of constituent materials and components in the top-level robot. Where possible, conducting all evaluations in the same simulator guarantees interoperability; nearly all simulators allow various materials properties to be defined and directly specified from lower levels. Multiscale modelling can enhance veracity. Some properties, for example hyperelasticity, can be tricky to model and may require specialist tools. In this case, co-simulation can be used to link specialist simulators together, allowing materials to be simulated, and their results shared with a dedicated component simulator to determine overall performance. Such approaches integrate with techniques that automatically validate the material models for use in simulators⁵⁶. Directly representing complex micro-level behaviours in higher-level models/simulators is difficult, but increasingly feasible as it receives ongoing research attention in multiple fields in materials science.

A key issue is the *reality gap*, where necessary abstractions lead to degraded performance when simulator-designed artefacts are transferred to reality. MLE heavily exploits techniques to reduce the gap. Gathering data on real designs and using a learning algorithm (e.g., a neural network⁵⁷ or Gaussian processes⁵⁸) to create surrogate models of the performance⁵⁹ improves accuracy and speed, and has been successful at the material level⁶⁰, the design level⁶¹, and for robot controllers⁶². Physical testing can improve an existing simulator to more closely match reality, either tuning simulator parameters^{63,64} and/or combining the predictions of the simulator with those of a data-driven model⁶⁵. In between these two ideas, it is possible to learn a ‘transferability function’ that predicts the accuracy of the simulator for a given design⁶⁶.

The Benefits of MLE

MLE is primarily designed to harness materials to provide diverse, specialised robot designs. Other main benefits include:

- **Scalability:** Promoting scalability is necessary due to combinatorial issues⁶⁷ brought about by embedding multiple search processes across three levels. Distributing the ‘full genome’ of a robot across multiple independent automatic design processes allows each level to be searched in parallel, using the specialist tools of each field where applicable to improve efficiency. Hybridisation shifts the majority of the search effort into relatively cheap, parallelisable simulations and models.
- **Self-optimisation:** Although the early stages of MLE are likely to be slow, with few options available, we envisage the system as somewhat self-optimising; the longer it runs, the better our models become, and the more options are discovered in every layer.
- **Re-use:** Focusing on physical properties allows materials and components to transfer between MLE architectures. Processes can be swapped in or out of a level with relative ease.
- **Collaboration:** MLE architectures will likely be distributed across multiple institutions depending on the availability of hardware and specialists, leading to an inherently collaborative effort integrating multiple research groups and the architecture itself, that promotes standardised, readily available experimental information and cross-fertilisation of ideas⁶⁸. We may look to the Materials Genome Initiative for inspiration on standardising MLE, encouraging collaboration, and reducing barriers to entry⁶⁹.

Opportunities for MLE architectures

As a new paradigm for designing robots, MLE will naturally gravitate towards certain applications. Consider the rapidly-advancing field of soft robotics^{70,71}; compliant, deformable robots that survive crushing, burning, and other hazards which are characteristic of the unstructured environments we want to put robots into. Integration of sensing, actuation, and deformation are fundamental to soft robotics; MLE most simply permits this using a single library of multi-function components, rather than dedicated sensing, actuation, etc.

Soft robotics currently lacks a codified design methodology⁷², as deformable soft materials are not amenable to conventional approaches⁷³. Designers often settle on a (frequently bio-inspired) preconceived design, for example an octopus or a jellyfish⁷⁴, which instantly places heavy limitations on the designs considered. Rather than design a fish, MLE lets us ask a different question - what ‘creatures’ might evolution devise if its building blocks were not protein, muscle, and bone; but rather polymers and composites? MLE is a perfect fit for the role of soft robot designer, harnessing diversity to comprehensively explore soft robot design spaces.

Soft robots have particularly interesting and powerful embodiments, which emerge through interacting arrangements of morphological and material properties^{75,76}. MLE provides a continuous stream of new materials and increasingly capable componentry^{77,78}, offering a pathway towards *designing for embodiment*: discovering specialized soft materials, fully leveraging those materials through the emergent generation of components and bodies, and creating controllers to strongly couple the resulting embodiments with the environment.

For our second design opportunity, let’s cast our minds forward 20-30 years. Imagine that we want to perform basic environment monitoring with robots: to traverse terrain in a zone, gather some data, and fully degrade after a while. This might sound simple at first, but critically depends on the environment. The Sahara is very hot, dry, and sunny (during the day), but Antarctica is cold and icy. Creepers and other low-lying foliage in the Amazon present a markedly different challenge to rolling desert sand dunes.

Designing robots for each niche with classic engineering would require an army of engineers for each environment, and the engineering cost would sky-rocket. This is why most of engineering is about standardization and not specialization. The alternative is MLE (Fig.2), which could automatically design suitable robots (unique combinations of materials, morphology, and behaviour) for each environment. They might resemble insects: relatively simple, small, highly integrated, highly specialized, and fit for function. Note that the same MLE architecture, with shared materials and identical objectives, could adapt robots to account for seasonal differences within a biome, or could design for each of the following environments, and provide the following features:

- Antarctica: wind-powered, sliding locomotion, water-resistant, degrades with heat (in the summer)
- The Amazon: crawling locomotion, degrades with humidity, bio-mass powered
- Sahara: solar-powered, sliding locomotion, heat resistant, degrades with UV

Towards a new era of Embodied Intelligence

The convergence in materials, manufacturing, and design paves the way towards radically new ways of producing robots. The main thesis of this Perspective is that MLE architectures can integrate different technologies and levels under an evolutionary umbrella. Considering that natural evolution succeeded in filling practically all environmental niches on Earth with highly adapted lifeforms, this approach holds great promise as a robotic




Fig. 2 | Showing how MLE can provide a diversity of robots for a diversity of environmental niches. In this alternate MLE architecture, each level consists of only one heterogeneous library of solutions. Such architectures are likely to be possible in the further future, where boundaries between sensing, actuation, and structure are collapsed to promote emergence and integration, at the cost of increased computational costs and combinatorial effort.

design technique. MLE is admittedly ambitious, and as such we have identified four key challenges to be overcome during an MLE research program.

- Initial designs will be constrained to materials that are easily created, characterised, and modelled. MLE materials search, together with high-throughput efforts globally, and advances in materials modelling will gradually alleviate this issue.
- Efficiency: In an attempt to create emergent embodiments by filling as many bins as possible, our conceptual prototype trades diversity for efficiency. To counter, imagine an ‘overseer’ program that greedily searches for good embodiments, allowing pointers to *any* material properties rather than experimentally confirmed or modelled properties, and subsequently skewing the search process to find materials with those properties if a promising embodiment is found. This leads to another issue around balancing the tension between any top-down influence and a bottom-up robot design process.
- Allocating resources across levels, e.g., including enough physical experimentation to keep simulations approximate to reality, which may be quantified and managed using Gaussian Processes to identify areas of uncertainty in our models. Bottlenecks are another issue; insufficient resource allocation to e.g., actuators may limit the range of final robot design. Using discretised libraries lets us estimate coverage, and allocate more resources to searches lacking coverage.
- Ideally, MLE would be fully autonomous. However the human designer will play a significant role for the foreseeable future; setting up (designing suitable measures of robot fitness, suitably discretising libraries, identifying suitable materials...) and running experiments (characterising materials, assembling and evaluating robots). Ongoing developments in automated characterisation, testing, and construction facilities will reduce this burden.

As well as challenges, we wish to highlight a significant opportunity; *representations* (recall that representations are mappings from genotype to phenotype). The main historical event in ER

was a move from direct (1-to-1 mappings) to indirect representations, a response to increased phenotypic complexity required in real-world artefacts. Hierarchical representation has received scant consideration to date; we see a huge opportunity for intelligent level-spanning representations, to describe complex artefacts *that are built from other artefacts* and their interactions.

Looking to the future, we envision a staged development of MLE systems:

- Stage 1. Within 5 years, prototype MLE systems will come online, spread across multiple research institutions. They will produce evolved robots in controlled lab settings, with heavy human intervention.
- Stage 2. In about a decade, MLE systems will be able to generate robots for a suitably constrained real-world mission. Increasingly integrated construction techniques will speed up evolution and reduce the amount of human intervention required.
- Stage 3, after around 20 years, will see deployments in real-world environments. As models become more sophisticated, and computing power more available, monolithic search processes will begin to merge, heightening the interplay between material and morphology and encouraging emergence (e.g., Fig 2).

Somewhat counter-intuitively for an architecture based on segregated levels, MLE is about collapsing boundaries; between research institutions, between scientific disciplines, between reality and virtuality, and between robots and their constituent materials. In doing so, we hope to create a holistic design process for a new type of robot, specialised all the way from material to machine.

Competing Interests

The authors declare no competing financial and non-financial interest.

Acknowledgements

DH, DFK, PV and DW would like to acknowledge Active Integrated Matter, one of CSIRO’s Future Science Platforms, for funding this research. JBM is funded by the European Union’s Horizon 2020 research and innovation programme (grant agreement number 637972, project “ResiBots”).

References

1. J. Carlson and R. R. Murphy. “How UGVs physically fail in the field”. In: *IEEE Transactions on robotics* 21.3 (2005), pp. 423–437.
2. C. G. Atkeson, P. W. Babu Benezun, N. Banerjee, D. Berenson, C. P. Bove, X. Cui, M. Dedonato, R. Du, S. Feng, P. Franklin, M. Gennert, J. P. Graff, P. He, A. Jaeger, J. Kim, K. Knoedler, L. Li, C. Liu, X. Long, and X. Xinjilefu. “What Happened at the DARPA Robotics Challenge Finals”. In: *The DARPA Robotics Challenge Finals: Humanoid Robots To The Rescue*. Springer Tracts in Advanced Robotics. 2018, pp. 667–684. ISBN: 978-3-319-74665-4.
3. L. Barrett. *Beyond the brain: How body and environment shape animal and human minds. A compelling description of the underlying principles of Embodied Cognition*. Princeton University Press, 2011.
4. R. Pfeifer and J. Bongard. *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.
5. S. B. Carroll, J. K. Grenier, and S. D. Weatherbee. *From DNA to diversity: molecular genetics and the evolution of animal design*. John Wiley & Sons, 2013.
6. R. A. Brooks. “Intelligence without representation”. In: *Artificial intelligence* 47.1-3 (1991), pp. 139–159.
7. T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever. “Evolution strategies as a scalable alternative to reinforcement learning”. In: *arXiv preprint arXiv:1703.03864* (2017).
8. T. C. Le and D. A. Winkler. “Discovery and Optimization of Materials Using Evolutionary Approaches”. English. In: *Chemical Reviews* 116 (May 2016), pp. 6107–6132. ISSN: 0009-2665. DOI: [10.1021/acs.chemrev.5b00691](https://doi.org/10.1021/acs.chemrev.5b00691).

9. G.-Z. Yang, P. Fischer, and B. Nelson. "New materials for next-generation robots". In: *Science Robotics* 2 (2017). DOI: [10.1126/scirobotics.aap9294](https://doi.org/10.1126/scirobotics.aap9294).
10. L. N. Soldatova, A. Clare, A. Sparkes, and R. D. King. "An ontology for a Robot Scientist". English. In: *Bioinformatics* 22 (July 2006), E464–E471. ISSN: 1367-4803. DOI: [10.1093/bioinformatics/btl207](https://doi.org/10.1093/bioinformatics/btl207).
11. W. F. Maier, K. Stoeve, and S. Sieg. "Combinatorial and high-throughput materials science". In: *Angewandte Chemie-International Edition* 46 (2007), pp. 6016–6067. ISSN: 1433-7851. DOI: [10.1002/anie.200603675](https://doi.org/10.1002/anie.200603675).
12. V. Sans and L. Cronin. "Towards dial-a-molecule by integrating continuous flow, analytics and self-optimisation". In: *Chemical Society Reviews* 45 (2016), pp. 2032–2043.
13. R. D. King. "Automating Chemistry and Biology Using Robot Scientists". English. In: *Ki 2015: Advances in Artificial Intelligence* 9324 (2015), pp. Xiv–Xv. ISSN: 0302-9743.
14. J. M. Granda, L. Donina, V. Dragone, D.-L. Long, and L. Cronin. "Controlling an organic synthesis robot with machine learning to search for new reactivity". In: *Nature* 559.7714 (July 2018). A recent, highly efficient robotic system that rapidly searches a space of materials through a combination of physical experimentation and machine learning., pp. 377–381. ISSN: 1476-4687. DOI: [10.1038/s41586-018-0307-8](https://doi.org/10.1038/s41586-018-0307-8).
15. S. Curtarolo, G. L. Hart, M. B. Nardelli, N. Mingo, S. Sanvito, and O. Levy. "The high-throughput highway to computational materials design". In: *Nature materials* 12.3 (2013), p. 191.
16. E. O. Pyzer-Knapp, C. Suh, R. Gomez-Bombarelli, J. Aguilera-Iparraguirre, and A. Aspuru-Guzik. "What Is High-Throughput Virtual Screening? A Perspective from Organic Materials Discovery". In: *Annual Review of Materials Research, Vol 45*. Ed. by D. R. Clarke. Vol. 45. Annual Review of Materials Research. 2015, pp. 195–216. ISBN: 978-0-8243-1745-4.
17. Y. Mengüç, N. Correll, R. Kramer, and J. Paik. "Will robots be bodies with brains or brains with bodies?" In: *Science Robotics* 2.12 (2017), eaar4527.
18. F. Calignano, D. Manfredi, E. P. Ambrosio, S. Biamino, M. Lombardi, E. Atzeni, A. Salmi, P. Minetola, L. Iuliano, and P. Fino. "Overview on additive manufacturing technologies". In: *Proceedings of the IEEE* 105 (2017), pp. 593–612. ISSN: 0018-9219.
19. L. Li, A. Haghighi, and Y. Yang. "A novel 6-axis hybrid additive-subtractive manufacturing process: Design and case studies". In: *Journal of Manufacturing Processes* 33 (2018), pp. 150–160.
20. P. Euijn, L. Giselle Hsiang, H. David, A. Henrique de Amorim, V. Mario Domingo Monzón, and P. Rubén. "A study of 4D printing and functionally graded additive manufacturing". In: *Assembly Automation* 37 (Apr. 2017), pp. 147–153. ISSN: 0144-5154. DOI: [10.1108/AA-01-2017-012](https://doi.org/10.1108/AA-01-2017-012).
21. J. Martínez, S. Hornus, H. Song, and S. Lefebvre. "Polyhedral Voronoi diagrams for additive manufacturing". In: *ACM Transactions on Graphics. Proceedings of SIGGRAPH 2018* 37.4 (Aug. 2018), p. 15. DOI: [10.1145/3197517.3201343](https://doi.org/10.1145/3197517.3201343).
22. T. Chen, J. Mueller, and K. Shea. "Integrated design and simulation of tunable, multi-state structures fabricated monolithically with multi-material 3D printing". In: *Scientific reports* 7 (2017), p. 45671.
23. G. Haghighi, E. Habtour, S.-H. Park, F. Gardea, and M. C. McAlpine. "3D printed electrically-driven soft actuators". In: *Extreme Mechanics Letters* 21 (2018), pp. 1–8. ISSN: 2352-4316. DOI: <https://doi.org/10.1016/j.eml.2018.02.002>.
24. A. Rosendo, M. Von Atzigen, and F. Iida. "The trade-off between morphology and control in the co-optimized design of robots". In: *PLoS one* 12.10 (2017), e0186107.
25. L. Brodbeck, S. Hauser, and F. Iida. "Morphological evolution of physical robots through model-free phenotype development". In: *PLoS one* 10.6 (2015). Describes a complete proof-of-concept generate-and-test system for evolving real-world robots., e0128444.
26. M. Wehner, R. L. Truby, D. J. Fitzgerald, B. Mosadegh, G. M. Whitesides, J. A. Lewis, and R. J. Wood. "An integrated design and fabrication strategy for entirely soft, autonomous robots". In: *Nature* 536 (2016). A prototype system that creates autonomous robots using a single self-contained fabrication process, pp. 451–455. ISSN: 0028-0836. DOI: [10.1038/nature19100](https://doi.org/10.1038/nature19100).
27. F. Silva, M. Duarte, L. Correia, S. M. Oliveira, and A. L. Christensen. "Open Issues in Evolutionary Robotics". In: *Evolutionary Computation* 24.2 (2016), pp. 205–236. DOI: [10.1162/EVCO_a_00172](https://doi.org/10.1162/EVCO_a_00172).
28. G. S. Hornby, J. D. Lohn, and D. S. Linden. "Computer-automated evolution of an X-band antenna for NASA's space technology 5 mission". In: *Evolutionary computation* 19.1 (2011), pp. 1–23.
29. J. E. Auerbach and J. C. Bongard. "Environmental influence on the evolution of morphological complexity in machines". In: *PLoS computational biology* 10.1 (2014). A powerful example of the application of artificial evolution to adapt robots to their environments., e1003399.
30. J. Lehman, J. Clune, D. Misevic, C. Adami, J. Beaulieu, P. J. Bentley, S. Bernard, G. Belson, D. M. Bryson, N. Cheney, et al. "The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities". In: *arXiv preprint arXiv:1803.03453* (2018).
31. S. Nolfi and D. Floreano. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press, 2000.
32. H. Lipson and J. B. Pollack. "Automatic design and manufacture of robotic lifeforms". In: *Nature* 406.6799 (2000), p. 974.
33. A. E. Eiben, S. Kernbach, and E. Haasdijk. "Embodied Artificial Evolution – Artificial Evolutionary Systems in the 21st Century". In: *Evolutionary Intelligence* 5.4 (2012), pp. 261–272.
34. A. E. Eiben and J. Smith. "From evolutionary computation to the evolution of things". In: *Nature* 521.7553 (2015). Highlights a 'new era' in using evolutionary algorithms to design physical entities (including robots), p. 476.
35. J. Rieffel, J.-B. Mouret, N. Bredeche, and E. Haasdijk. "Introduction to the evolution of physical systems special issue". In: *Artificial Life* (2017).
36. N. Cheney, R. MacCurdy, J. Clune, and H. Lipson. "Unshackling Evolution: Evolving Soft Robots with Multiple Materials and a Powerful Generative Encoding". In: *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation. GECCO '13*. New York, NY, USA: ACM, 2013, pp. 167–174. ISBN: 978-1-4503-1963-8. DOI: [10.1145/2463372.2463404](https://doi.org/10.1145/2463372.2463404).
37. A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computation*. Springer-Verlag, 2003. ISBN: 3-540-40184-9.
38. K. O. Stanley. "Compositional pattern producing networks: A novel abstraction of development". In: *Genetic programming and evolvable machines* 8.2 (2007), pp. 131–162.
39. S. Doncieux, N. Bredeche, J.-B. Mouret, and A. E. G. Eiben. "Evolutionary Robotics: What, Why, and Where to". In: *Frontiers in Robotics and AI* 2 (2015), p. 4. ISSN: 2296-9144. DOI: [10.3389/frobt.2015.00004](https://doi.org/10.3389/frobt.2015.00004).
40. H. Rabitz. "Control in the sciences over vast length and time scales". In: *Quantum Physics Letters* 1.1 (2012), pp. 1–19.
41. C. Hansch, A. Leo, D. Hoekman, et al. *Exploring QSAR: fundamentals and applications in chemistry and biology*. Vol. 557. American Chemical Society Washington, DC, 1995.
42. J.-B. Mouret and J. Clune. "Illuminating search spaces by mapping elites". In: *arXiv preprint arXiv:1504.04909* (2015).
43. J. K. Pugh, L. B. Soros, and K. O. Stanley. "Quality diversity: A new frontier for evolutionary computation". In: *Frontiers in Robotics and AI* 3 (2016). Describes evolutionary algorithms whose success hinges on diversity., p. 40.
44. V. Vassiliades, K. Chatzilygeroudis, and J.-B. Mouret. "Using Centroidal Voronoi Tessellations to Scale Up the Multidimensional Archive of Phenotypic Elites Algorithm". In: *IEEE Transactions on Evolutionary Computation* 22.4 (2018), pp. 623–630.
45. K. S. et al. "Designing Neural Networks through Neuroevolution". In: *Nature Machine Intelligence* xxxx (2019), To be published this issue.
46. M. A. Figueiredo et al. "Adaptive sparseness for supervised learning". In: *IEEE transactions on pattern analysis and machine intelligence* 25.9 (2003), pp. 1150–1159.
47. R. Tibshirani. "Regression shrinkage and selection via the lasso". In: *Journal of the Royal Statistical Society. Series B (Methodological)* (1996), pp. 267–288.
48. A. Cully and Y. Demiris. "Hierarchical Behavioral Repertoires with Unsupervised Descriptors". In: *Genetic and Evolutionary Computation Conference (GECCO)*. ACM, 2018.
49. S. Kriegman, N. Cheney, F. Corucci, and J. C. Bongard. "Interceptive robustness through environment-mediated morphological development". In: *arXiv preprint arXiv:1804.02257* (2018).
50. H. Hauser, A. J. Ijspeert, R. M. Fuchslin, R. Pfeifer, and W. Maass. "Towards a theoretical foundation for morphological computation with compliant bodies". In: *Biological cybernetics* 105.5-6 (2011), pp. 355–370.
51. A. E. Eiben, N. Bredeche, M. Hoogendoorn, J. Stradner, J. Timmis, A. Tyrrell, and A. F. T. Winfield. "The Triangle of Life: Evolving Robots in Real-time and Real-space". In: *Advances in Artificial Life, ECAL 2013*. Ed. by P. Liò, O. Miglino, G. Nicosia, S. Nolfi, and M. Pavone. Taormina, Italy: MIT Press, 2013, pp. 1056–1063. ISBN: 9780262317092. DOI: [10.7551/978-0-262-31709-2-ch157](https://doi.org/10.7551/978-0-262-31709-2-ch157).
52. A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret. "Robots that can adapt like animals". In: *Nature* 521.7553 (2015), p. 503.
53. M. Bhattacharya. "Evolutionary approaches to expensive optimisation". In: *arXiv preprint arXiv:1303.2745* (2013).

54. D. Howard. "A Platform that Directly Evolves Multirotor Controllers". In: *IEEE Transactions on Evolutionary Computation* (2017). ISSN: 1089-778X.
55. H. Heijnen, D. Howard, and N. Kottege. "A testbed that evolves hexapod controllers in hardware". In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. 2017, pp. 1065–1071. DOI: [10.1109/ICRA.2017.7989128](https://doi.org/10.1109/ICRA.2017.7989128).
56. T. Bäck, L. Keßler, and I. Heinle. "Evolutionary strategies for identification and validation of material model parameters for forming simulations". In: *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. ACM. 2011, pp. 1779–1786.
57. M. Hüsken, Y. Jin, and B. Sendhoff. "Structure optimization of neural networks for evolutionary design optimization". In: *Soft Computing* 9.1 (2005), pp. 21–28.
58. C. K. I. Williams and C. E. Rasmussen. "Gaussian processes for machine learning". In: *the MIT Press* 2.3 (2006), p. 4.
59. Y. Jin. "Surrogate-assisted evolutionary computation: Recent advances and future challenges". In: *Swarm and Evolutionary Computation* 1.2 (2011), pp. 61–70.
60. D. A. Winkler and T. C. Le. "Performance of deep and shallow neural networks, the universal approximation theorem, activity cliffs, and QSAR". In: *Molecular informatics* 36.1-2 (2017), p. 1600118.
61. A. Gaier, A. Asteroth, and J.-B. Mouret. "Data-Efficient Design Exploration through Surrogate-Assisted Illumination". In: *Evolutionary Computation* (2018), pp. 1–30.
62. K. Chatzilygeroudis, R. Rama, R. Kaushik, D. Goepp, V. Vassiliades, and J.-B. Mouret. "Black-box data-efficient policy search for robotics". In: *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on*. IEEE. 2017, pp. 51–58.
63. J. Bongard, V. Zykov, and H. Lipson. "Resilient machines through continuous self-modeling". In: *Science* 314.5802 (2006), pp. 1118–1121.
64. J. C. Zagal and J. Ruiz-Del-Solar. "Combining simulation and reality in evolutionary robotics". In: *Journal of Intelligent and Robotic Systems* 50.1 (2007), pp. 19–39.
65. K. Chatzilygeroudis and J.-B. Mouret. "Using Parameterized Black-Box Priors to Scale Up Model-Based Policy Search for Robotics". In: *IEEE International Conference on Robotics and Automation*. 2018.
66. S. Koos, J.-B. Mouret, and S. Doncieux. "The transferability approach: Crossing the reality gap in evolutionary robotics". In: *IEEE Transactions on Evolutionary Computation* 17.1 (2013), pp. 122–145.
67. D. L. Donoho. "High-dimensional data analysis: The curses and blessings of dimensionality". In: *AMS Math Challenges Lecture 1* (2000), p. 32.
68. M. D. Wagdy and J. C. Bongard. "Combining computational and social effort for collaborative problem solving". In: *PloS one* 10.11 (2015), e0142524.
69. A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, et al. "Commentary: The Materials Project: A materials genome approach to accelerating materials innovation". In: *Apl Materials* 1.1 (2013), p. 011002.
70. M. T. Tolley, R. F. Shepherd, B. Mosadegh, K. C. Galloway, M. Wehner, M. Karpelson, R. J. Wood, and G. M. Whitesides. "A resilient, untethered soft robot". In: *Soft Robotics* 1 (2014), pp. 213–223.
71. D. Rus and M. T. Tolley. "Design, fabrication and control of soft robots". In: *Nature* 521 (2015), pp. 467–475. ISSN: 0028-0836.
72. H. Lipson. "Challenges and opportunities for design, simulation, and fabrication of soft robots". In: *Soft Robotics* 1.1 (2014), pp. 21–27.
73. J. C. Bongard. *Evolving Soft Robots*. 2016.
74. S.-W. Yeom and I.-K. Oh. "A biomimetic jellyfish robot based on ionic polymer metal composite actuators". In: *Smart materials and structures* 18 (2009), p. 085002. ISSN: 0964-1726.
75. M. Manti, V. Cacucciolo, and M. Cianchetti. "Stiffening in Soft Robotics A Review of the State of the Art". English. In: *IEEE Robotics & Automation Magazine* 23 (Sept. 2016), pp. 93–106. ISSN: 1070-9932. DOI: [10.1109/Mra.2016.2582718](https://doi.org/10.1109/Mra.2016.2582718).
76. S. Bauer, S. Bauer-Gogonea, I. Graz, M. Kaltenbrunner, C. Keplinger, and R. Schwödiauer. "25th anniversary article: a soft future: from robots and sensor skin to energy harvesters". In: *Advanced Materials* 26.1 (2014), pp. 149–162.
77. S.-T. Han, H. Peng, Q. Sun, S. Venkatesh, K.-S. Chung, S. C. Lau, Y. Zhou, and V. A. L. Roy. "An Overview of the Development of Flexible Sensors". In: *Advanced Materials* (2017). ISSN: 1521-4095.
78. A. Miriyev, K. Stack, and H. Lipson. "Soft material for soft actuators". In: *Nature Communications* 8 (2017).