

t for Two

Linear Synergy Advances the Evolution of Directional Pointing Behaviour

Authors

Institute

Institute

email

Abstract. Motor synergies, i.e. systematic relations between effectors, have been first proposed as a principle in motor control by N. Bernstein in 1935. Thereafter, his idea has inspired many models of motor control in humans and animals. Recently, “linear synergy”, i.e. a linear relation between the torques applied to different joints, was reported to occur in human subjects during directional pointing movements [4]. In this paper, results from experiments in evolutionary robotics are presented to explore the concept of synergies in general and the role of linear synergy in the organisation of movement in particular. A 3D simulated robotic arm is evolved to reach to different target spots on a plane. Linear synergy is not found to be an outcome of the evolutionary search process, but imposing linear synergy as a constraint on artificial evolution dramatically improves evolvability and performance of evolved controllers.

1 Motor Synergies

*The centipede was happy quite,
Until the toad in fun
Said “Pray, which leg goes after which?”
Which worked his mind to such a pitch,
He lay distracted in a ditch,
Considering how to run.*

This anonymous ditty expresses nicely what the Russian physiologist N. Bernstein, a pioneer in biomechanics and anticipator of ideas of cybernetics, has identified as the degrees of freedom (DoF) problem as early as 1935 (English translation 1967 [3]): if we conceive of the central nervous system (CNS) as a homuncular control organ that determines the state of all actuators at any point in time, the control problem it has to solve is of inconceivably high dimensionality. Describing human or animal motion in terms of joint kinematics already involves a large number of DoFs (e.g. 7 in moving an arm), but if motor control is thought of in terms of individual muscles, or even motor neurons, the number of DoFs to be controlled just for moving an arm quickly exceeds four digits [12].

Not only is the problem space intractably large, but it is also *redundant* with respect to the outcome of an action, a condition that Hebb has termed “Motor

Equivalence” ([6], p. 153ff). A famous example for motor equivalence is the fact that “handwriting appears about the same regardless of the size of the letters or the limb or body segment used to produce them” ([7], p.659). Yet another difficulty in practising motor control is *context conditioned variability* ([12], p. 246ff): the effect of a motor command is sensitive to the anatomical, mechanical and physiological context of the interaction of an agent with its environment, e.g. limb positions, passive dynamics or the state of the peripheral nervous system.

Motor synergies, i.e. *systematic relations between effectors*, are Bernstein’s solution to the DoF problem. Just as the driver of a car, due to the linkage of the two front wheels, can determine both their positions by using just one steering wheel, mutual constraints in an organism’s motor system could serve to build functional subunits, thereby reducing the effective number of DoFs in a motor task. Although this idea of motor synergies has greatly influenced research in motor behaviour (e.g. [1, 5, 8, 10]), it is not free from practical and conceptual problems: is explaining the CNS as the driver of a bodily car much easier than explaining the whole system? And, even more inscrutable, where do synergies come from, what is a good synergy, what mechanism controls their development and maintenance and the interaction between synergies acting in parallel? Does the evidence that “the context in which a motor task is executed strongly influences its organization” ([14], p. 74) not contradict the idea of low level synergies organising movement primitives in a constrained and automated manner?

Situated and embodied approaches to behaviour and cognition offer a way around the problems outlined: typically, researchers working within these paradigms do not conceive of motion as a principally feed-forward process, planned by a homunculus in a disembodied control chamber, to later be interpreted by “low level” mechanisms in terms of bodily coordinates. In such non-homuncular explanations, it is unclear what functional role motor synergies could bear. However, motor synergies are evident in human and animal behaviour (for a summary of findings see e.g. [13]). It is exactly the puzzle of their purpose that make synergies an interesting phenomenon to explore from a situated and embodied perspective, in particular within an evolutionary robotics framework. The following are some of the questions that can be addressed: Where do synergies come from? Under which circumstances do they arise? Are synergies just epiphenomenal to a structured agent environment interaction or implemented in the control architecture of an evolved agents? In the reverse direction, thinking in terms of synergies can also enrich our understanding and practice of evolving artificial agents: in the endeavour of staying tractable, many evolutionary robotics experiments seem to *implicitly presuppose the existence of synergies*: robot control architectures are frequently not redundant in DoFs and the role of actuators is rarely subject to context conditioned variability, as they are in biological organisms. Any comparison of the processes realising a certain behaviour in a robot and those realising the same behaviour in a human or animal will always rely on synergy-like low level processes to compensate for all such idealisations. To study

the impact of idealising assumptions can help to improve our use of evolutionary robotics as a method for the study of intelligent behaviour.

In this paper, the role of motor synergies and redundancy in DoFs in motor control is investigated by performing a systematic comparison of evolvability and evolved behaviour in different kinds of controllers. The directional pointing task chosen follows empirical experiments where linear synergies were observed in human subjects. If linear synergies fulfil a purpose in the organisation of the modelled task, their existence should be reflected by an improvement in either performance or evolvability. The results are analysed with respect to methodological issues in evolutionary robotics and put into relation with empirical findings.

2 Experiments in Directional Pointing

The experiments presented in this paper are inspired by findings reported in Gottlieb et Al. [4] on the linear relation between the torque applied to the shoulder and the torque applied to the elbow during human directional pointing in the sagittal plane. This principle of linear synergy could not be expected from the nature of the task and does not appear to be the outcome of a learning process, since infants in the pre-reaching period already apply it, even though their attempts to grasp an object are unsuccessful [15]. The role linear synergy plays in the realisation of pointing movements remains mysterious.

Following these empirical findings, a simulated robotic arm is evolved to reach to six different targets in the horizontal plane. The arm is simulated in the C++ open source physics simulation library Open Dynamics Engine (ODE, [11]). It consists of a forearm, an upper arm (each two units long) and a spherical hand (Fig. 1). The six target points are spread evenly with uniformly distributed noise $\epsilon [0, \frac{1}{6}\pi]$ on the circumference of a circle with a radius of 1.25 units around the unitary starting position of the hand at an elbow angle of 60° (Fig. 1, inlay).




Fig. 1. The simulated robotic arm and its DoFs. Inlay: Plan view of the task.

Every DoF in the model is controlled by applying a torque t_i to a joint j_i . In order to test the effect that the number of DoFs has on the task, experiments are run on a planar (**2D**) condition, with one DoF in the elbow joint (t_E) and one in the shoulder (t_{S1}), and a three dimensional (**3D**) condition, with one DoF in the elbow (t_E) and three DoFs in the shoulder (rotation in the horizontal plane (t_{S1}), lifting/lowering (t_{S2}) and rotation about the arm length (t_{S3}), Fig. 1). Joint stops are applied following the human example. In order to keep the task complexity comparable in the **2D** and the **3D** condition, deviation of the hand from the horizontal plane in the **3D** condition is made impossible. Thereby, the

model loses biological plausibility, but not the suitability to explore the role of redundant DoFs in a motor task in principle. Dry friction is applied at all joints, gravity is not modelled.

The control networks evolved are continuous time recurrent neural networks (CTRNNs, e.g.[2]) with an input layer projecting to a fully connected hidden layer, which again projects in a feed forward fashion to the output neurons. The dynamics of neurons n_i in a CTRNN of N neurons are governed by

$$\tau_i \frac{da_i(t)}{dt} = -a_i(t) + \sum_{j=0}^N w_{ij} \sigma(a_j(t) + b_j) + I_i \quad (1)$$

where $\sigma(x) = \frac{1}{1+e^{-x}}$ is the standard sigmoidal function and I_i is the external input to n_i . The weights $w_{ij} \in [-7, 7]$ from n_j to n_i , the bias $b_i \in [-3, 3]$ and the time constant $\tau_i \in [0.01, 1.77]$ are determined by a genetic algorithm (GA).

Three different conditions are investigated for both the **2D** and **3D** set-ups.

In the unconstrained (**UC**) condition, torques t_i are generated by $t_i = MG_i \cdot (a_x - a_y)$ where $MG_i \in [0.1, 30]$ is the evolved motor gain and $a_{x,y}$ is the activity of the two antagonistic motor neurons dedicated to the control of the joint j_i . The network has sensory neurons giving the angular position of each DoF, one sensory neuron for the pointing direction $\theta \in [0, 2 \cdot \pi]$ and six hidden neurons.

In the ‘‘split brain’’ (**SB**) condition, two different controllers for the different joints are co-evolved following the architecture of the **UC** condition. They share the directional input neuron, but are only fed the corresponding joint angles j_i and have just three hidden neurons each. Comparing results from the **SB** and the **UC** condition is interesting with respect to role of a *neural basis for synergies*.

Finally, in the forced synergy (**FS**) condition, only the torque t_E to the elbow is generated by a CTRNN, the other joint torques t_{Sj} are scaled as a linear function $t_{Sj} = K_j \cdot t_E$ where K_j is constant within a pointing movement, but varies systematically across trials with the desired pointing direction: $K_j = f(\theta)$. Two different functional representations are used, in a condition called **FSa**, $K(\theta)$ is a simple linear function

$$K_j^a(\theta) = k_j^1 \cdot \theta + k_j^2 \quad (2)$$

with $k_j^i \in [-4, 4]$ set genetically. In the condition **FSb**, $K_j(\theta)$ is represented by a radial basis function (RBF) network with Gaussian RBFs

$$K_j^b(\theta) = \sum_{i=1}^4 w_{Ri} \cdot e^{-\frac{\delta^2}{2 \cdot \Delta^2}} \quad (3)$$

where $\delta = c_i - \theta$, $d \in [-\pi, \pi]$ is the difference in direction between the evolved RBF center $c_i \in [-\pi, \pi]$ and the target direction θ . The width of the Gaussian RBF $\Delta \in [0.5, 1.5]$ and the weights $w_{Ri} \in [-4, 4]$ are set genetically. The **FS** networks do not receive θ as neuronal input.

All sensory inputs in all conditions are multiplied by a genetically set sensor gain $SG_i \in [0.1, 20]$. The behaviour of the network is simulated using Euler integration with a time-step of 0.01 time units, the same time step as used in the simulation of the arm in ODE. Trials are run for $T \in [2000, 3000]$ time steps.

The parameters of the control networks are evolved in a population of 30 individuals with a generational genetic algorithm with real-valued genes $\epsilon \in [0, 1]$, truncation selection ($\frac{1}{3}$), vector mutation ([2]) of magnitude $r = 0.6$ and reflection at the gene boundaries. All values are mapped linearly to the target range, apart from the sensor gains SG_i , the motor gains MG_i , the time constants τ_i , the absolute values of the coefficients $|k_i|$ and the absolute values of the RBFN weights $|w_{Ri}|$, which are mapped exponentially.

The fitness $F(i)$ of an individual i on a target spot j is given by

$$F_j(i) = 1 - \frac{d_j(T, i)^2}{d_j(0, i)^2} \quad (4)$$

where $d_j(t, i)$ is the distance of the hand from the target spot j at time t for individual i . Networks for all conditions are evolved with either incremental evolution, where the next clockwise target spot is added to the evaluation once the average performance of the population exceeds $\bar{F} = 0.4$ (starting with two positions) or on all six target spots right from the start. The evaluation of a network i on n target spots is given by

$$F(i) = \sum_{j=1}^n F_j(i) \cdot 2^{-(j-1)} \cdot \frac{1}{\sum_{j=1}^n 2^{-(j-1)}} \quad (5)$$

where $F_j(i)$ gives the fitness on the j^{th} worst evaluation trial for individual i , which gives more weight to worse evaluations and thereby rewards the generalisation capacity of the evolved networks.

3 Results

3.1 Number of Degrees of Freedom

Even though the investigated pointing task is already redundant in the **2D** condition because of the infinite possible trajectories leading the hand to the target position, the liberation of movement by providing two additional DoFs in the **3D** condition results in an enormous increase in performance in all network architectures: the number of target spots reached is much higher (Fig. 3).

Although the temporal organisation of movement can vary in **2D** networks, they do not have other possibilities for solving the task than to bring the two planar joints in the appropriate end positions. The **3D** networks, contrariwise, make excessive use of the additional DoFs and exploit passive dynamics to exhibit a multitude of strategies for solving the pointing task. An invariant in the behaviour of the **UC** and the **SB** networks was that their solutions involved turning the arm along its length to one of the joint stops — apparently, the positions thereby reached are more suitable for evolutionary search than the original starting position.

The performance of evolved controllers i was tested by selectively “anaesthetising” DoFs, i.e. not applying the motor torques, but allowing passive dynamics ($F'_a(i)$) and by blocking them ($F'_b(i)$). Figure 2 shows the squared difference in performance $(F'_a(i) - F'_b(i))^2$ between those conditions: where enabling passive dynamics to work on the anaesthetised DoFs does hardly make a difference in the **2D** condition (Fig. 2 (A)), it has a noticeable impact on performance of all **3D** networks. Passive dynamics is a source of behavioural stability for controllers in the redundant **3D** condition.




Fig. 2. Squared difference in normalised performance as individual joints are free to move but not driven (F'_a) or blocked (F'_b) in exemplary **2D** (A) and **3D** (B) networks.

These findings illustrate how suppressing the exploitation of passive dynamics and endowing an agent with the minimally required sensorimotor system for a task can bias and limit evolved behaviours and even hamper evolution, even though dimensionality of the search space is reduced in the minimalist case.

3.2 Forcing Linear Synergy

Probably the most significant finding from the presented experiments is the dramatic advantage that the **FSa** and **FSb** networks have in performance and evolvability over both the **UC** and **SB** networks. Figure 3 shows how, in incremental evolution, the **FSb** networks advance to the next goal twice as many times as the non-**FS** networks. With twice as many generations, the non-**FS** networks come closer to, but never reach the level of performance of the **FS** networks. The **3D FSb** network is the only one that succeeds in solving the entire problem space, average performance of best individuals after 1000 generations is 0.65. Non-incremental evolution led to qualitatively similar results.

Where it could be argued that the RBFN is simply a very suitable representation for a scaling function in this task, this is certainly not the case for a simple linear function, particularly in the **2D** case, where the sensorimotor system of the robot is already so restricted that adding this additional harsh constraint makes it impossible to generate a controller that masters the task, not just because of the singularity at $\theta = 2\pi$. To rule out the possibility that




Fig. 3. Average number of starting positions reached in incremental evolution after 100 (dark) and 500 (light) generations across ten evolutionary runs.

the **UC** and **SB** controllers simply could not cope with the presentation of the input direction as a scalar neural input, a more “CTRNN friendly” set-up was tested, where controllers were provided with six different input neurons for the different target spots and no noise applied to θ , but neither in the **3D** or the **2D** condition could they go beyond three targets within 1000 generations.

Evolving controllers for the outlined directional pointing task under the constraint of linear synergy could be shown to significantly improve both evolvability and performance of the resulting controllers. Even when trying to remove all biases in the experimental set-up that could give the **FS** networks a task-specific advantage, this advantage was preserved. The division of control into scaling function and generation of motor signal is in some way very suitable for evolutionary search. The nature of this benefit, however, is unclear, as analysis of the ruggedness of the fitness landscape around successfully evolved individuals did not provide an explanation.

3.3 Evolved Synergies

As for the linear synergies $K_i(\theta)$ evolved, no general pattern could be observed. Figure 4 (A) depicts an exemplary evolved RBFN in the **3D FSb** condition: the systematicities with which the scaling constant $K_i(\theta)$ varies in the different DoFs do not stand in an obvious relation. The displacement and overlap of peaks in these functions explain the diversity of behavioural strategies for different domains of θ observed in the **FSb** networks: For different targets, different DoFs are predominant in the realisation of the task.

If the principle of linear synergy was a characteristic of a good solution to the pointing task, we would expect an increase of linearity in torque relation as the performance of the **UC** and **SB** networks increases. Figure 4 (B) shows the sum of squared error from linear synergy in the **2D** and **3D UC** and **SB** controllers in the best individuals across 5 evolutionary searches: Where a slight tendency to get closer to linear synergy as performance increases is noticeable in the **2D** networks, in the **3D** networks, linear synergy and performance appear to be completely unrelated. It is remarkable, though, that on average, the **SB** networks act much less in linear synergy (note logarithmic scale), even though variance in the **SB** networks is much higher. From the mere architecture of the




Fig. 4. (A): An exemplary evolved RBFN for a **3D FSb** network. (B): Sum of squared error from linear synergy across generations in the **2D** (top) and **3D** (bottom) networks, solid: **UC**, dashed: **SB**, average across five evolutions. (Note non-linear scales) (C): A **2D UC** network (bottom) applying a similar strategy as a **2D FSb** network (top).

controllers, this tendency to not simultaneously activate joints can be expected.¹ However, being more disposed to linear synergy does not seem to give the **UC** controllers an advantage over the **SB** controllers, suggesting that the magnitude of deviation from linear synergy is not essential. Fig. (C) shows, how a **2D UC** network applies a very similar strategy for solving a task, but slightly displaces the torque peaks, which leads to a loop in the t_E/t_{S1} map. On a performance level, this deviation from linear synergy does not mean a disadvantage.

In summary, a direct correlation between the level of performance and the level of linear synergy is not evident, particularly in the **3D** networks, where the deviation from linear synergy in the random initial populations is maintained throughout evolution. Hence, the principle of linear synergy is not an a priori good strategy for solving the pointing task — it is constraining the search space that implements the crucial advantage linear synergy offers for the evolution of directional pointing behaviour.

4 Discussion

Inspired by empirical evidence, a series of evolutionary robotics experiments have been conducted to shed light on the opaque role of linear synergy in a directional pointing task. Linear synergies could not be found to be the outcome of an unconstrained evolutionary search process. Furthermore, disconnected controllers for the different joints did not have a significant evolutionary disadvantage compared to monolithic networks controlling both joints, suggesting that the mere possibility of implementing constraints between effectors in a network does not provide a selective advantage. However, evolution under the constraint of linear synergy boosts the evolutionary search process, even with an impoverished linear scaling functions ($K_i^a(\theta)$), but more significantly with a RBFN representing

¹ It may be worth mentioning that the lack of neural connections between the controllers in the **SB** networks does not rule out the formation of synergies. Regularities in activation could in principle be mediated through the environment. Discovering such synergies would pose a clear challenge to homuncular explanations of synergies.

the systematic variation of the linear scaling constant. The benefits of passive dynamics and redundant DoFs became clear during analysis.

For the investigated task, both a complication (i.e. adding more DoFs) and a restriction (i.e. forcing linear synergy) of the space to be searched by a GA could be shown to provide an evolutionary advantage, independent of each other. Thus, improving evolvability is not a matter of scaling up or scaling down the search space, but of *reshaping the fitness landscape*. Once tasks and robotic platforms become more complex, evolutionary robotics will have to come up with according reshaping techniques to scaffold the evolutionary search process and thereby solve the “bootstrap problem” ([9], p. 13). Such a reshaping always means biasing the evolutionary search process. A lesson that can be learned from the results presented in this paper is maybe that introducing biases following a biological example can not help in not only making a stronger point when relating findings obtained to empirical findings from biology or ethology, but can also be beneficial from a plain engineering point of view.

A generalisation about the evolvability of the applied technique, however, is not justified. Apart from technical difficulties, such as defining where one movement starts and another one ends or formalising the context of the task as a set of variables according to which a scaling constant K_i can vary, in many tasks a linear relation between effectors will be disadvantageous: considering, for instance, a two-wheeled robot performing obstacle avoidance; it seems quite obvious that its success relies on a constant change in relation between the two effectors. A hypothesis put forward in this paper is that the principle of linear synergy will provide an evolutionary advantage in tasks that are not primarily reactive or if the motor system of a robot is redundant in DoFs. Otherwise, the imposed restrictions are in direct opposition to the required behaviour.

When asking the question of what the presented findings can teach us about the role of synergies in human motor behaviour, conclusions have to be drawn carefully. Neither the **SB** networks nor the **UC** networks could be observed to increase linear synergy as performance increases, so linear synergy does not seem to be a priori a good strategy to master a pointing task. The enormous advantage that imposed linear synergy has on the evolutionary search process, however, suggests that it may be a useful constraint for a developmental process. This result relates well with the findings reported by Zaal et Al. [15] about infants employing linear synergy regardless of movement success. The authors hypothesise that linear synergy eases the acquisition of reaching and pointing movements at an early stage, a hypothesis that is supported by the experimental results. In order to further investigate this hypothesis, it would be interesting to study the phylogeny of linear synergy, or, as an extension to the experiments presented here, to evolve the constraints for an ontogenetic developmental process. It would be desirable to abolish the restriction of hand movement to the plane in further experiments, in the spirit of being closer to the biological example.

Even though the original notion of motor synergies was embedded in a homuncular view of motor control which is generally rejected by dynamical and artificial life perspectives, the investigation of motor behaviour in terms of syn-

ergies is a fruitful way of thinking. Within an evolutionary robotics framework, systematicities between effectors, as they are ubiquitous in humans and animals, can be investigated in order to find out about the different functions they may bear. On the other hand, as robots and tasks become more complex, description and explanation of the behaviour obtained become more complex as well. Looking for systematicities between effectors can be a good starting point in trying to understand, and designing experiments informed by observations on biological organisms may help to generate intelligent behaviour.

References

1. Arbib, M. A.: *Perceptual Structures and Distributed Motor Control*. In: V. B. Brooks (ed.): *Handbook of Physiology. Section 2: The Nervous System*. Vol. II, Motor Control, Part 1. American Physiological Society (1981) 1449–1480.
2. Beer, R. D.: *Toward the Evolution of Dynamical Neural Networks for Minimally Cognitive Behavior*. In: P. Maes, M. J. Mataric, J.-A. Meyer, J. B. Pollack & S. W. Wilson (eds.): *From Animals to Animats 4. Proc. 4th Int. Conf. on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press 1996. 421–429.
3. Bernstein, N.: *The Coordination and Regulation of Movements*. Oxford: Pergamon 1967.
4. Gottlieb, G. L., Q. Song, G. L. Almeida, D. Hong, and D. Corcos: *Directional Control of Planar Human Arm Movement*. *Journal of Neurophysiology* 78 (1997) 2985–2998.
5. Grossberg, S. and Paine, R.W.: *A Neural Model of Corticocerebellar Interactions During Attentive Imitation and Predictive Learning of Sequential Handwriting Movements*. *Neural Networks*, 13 (2000) 999–1046.
6. Hebb, D.O.: *The Organization of Behavior. A Neuropsychological Theory*. John Wiley & sons inc., New York 1949.
7. Kandel, Eric R., James H. Schwartz and Thomas M. Jessel (eds.): *Principles of Neural Science*. Fourth Edition. McGraw–Hill, New York 2000.
8. Morasso, P., F.A. Mussa Ivaldi and C. Ruggiero: *How a Discontinuous Mechanism Can Produce Continuous Patterns in Trajectory Formation and Handwriting*. *Acta Psychologica* 54 (1983) 83–98.
9. Nolfi, S and D. Floreano: *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge MA 2000.
10. Sporns, O., and G.M. Edelman: *Solving Bernstein’s Problem: A Proposal for the Development of Coordinated Movement by Selection*. *Child Dev.* 64 (1993) 960–981.
11. Smith, R.: *Open Dynamics Engine. 0.5 release*. <http://ode.org>, retrieved: 10.1.2005.
12. Tuller, B., H. Fitch and M. Turvey: *The Bernstein Perspective: II. The Concept of Muscle Linkage or Coordinative Structure*. in: S. Kelso (ed.): *Human Motor Behavior. An Introduction*. Hillsdale: Lawrence Erlbaum (1982) 239–252.
13. Turvey, M., H. Fitch and B. Tuller: *The Bernstein Perspective: I. The Problems of Degrees of Freedom and Context-Conditioned Variability*. in: S. Kelso (ed.): *Human Motor Behavior. An Introduction*. Hillsdale: Lawrence Erlbaum (1982) 253–270.
14. Weiss, P. and M. Jeannerod: *Getting a Grasp on Coordination*. *News Physiol. Sci.* 13 (1998) 70–75.
15. Zaal, F., Daigle, K., Gottlieb, G.L., Thelen, E.: *An Unlearned Principle for Controlling Natural Movements*. *Journal of Neurophysiology*, 82 (1999) 255–259.