Defence Technology xxx (xxxx) xxx

Contents lists available at ScienceDirect

Defence Technology

journal homepage: www.keaipublishing.com/en/journals/defence-technology

Adaptive optimisation of explosive reactive armour for protection against kinetic energy and shaped charge threats

Philipp Moldtmann [a], Julian Berk [b], Shannon Ryan [b, *], Andreas Klavzar [c], Jerome Limido [d], Christopher Lange [a, c], Santu Rana [b], Svetha Venkatesh [b]

[a] University of the Federal Armed Forces Hamburg, Holstenhofweg 85, Hamburg 22043, Germany
[b] Applied Artificial Intelligence Institute (A^2I^2), Deakin University, 75 Pigdons Rd, Waurn Ponds, VIC 3216, Australia
[c] French-German Research Institute of Saint-Louis (ISL), 5 Rue de General Cassagnou, Saint-Louis 68300, France
[d] ABSTRAO, 16 Rue de Isaac Newton, Plaisance-du-Touch 31830, France

ARTICLE INFO

Article history:
Received 20 March 2024
Received in revised form
25 April 2024
Accepted 13 May 2024

Keywords:
Terminal ballistics
Armour
Explosive reactive armour
Optimisation
Bayesian optimisation

ABSTRACT

We evaluate an adaptive optimisation methodology, Bayesian optimisation (BO), for designing a minimum weight explosive reactive armour (ERA) for protection against a surrogate medium calibre kinetic energy (KE) long rod projectile and surrogate shaped charge (SC) warhead. We perform the optimisation using a conventional BO methodology and compare it with a conventional trial-and-error approach from a human expert. A third approach, utilising a novel human-machine teaming framework for BO is also evaluated. Data for the optimisation is generated using numerical simulations that are demonstrated to provide reasonable qualitative agreement with reference experiments. The human-machine teaming methodology is shown to identify the optimum ERA design in the fewest number of evaluations, outperforming both the stand-alone human and stand-alone BO methodologies. From a design space of almost 1800 configurations the human-machine teaming approach identifies the minimum weight ERA design in 10 samples.

1. Introduction

Explosive reactive armour (ERA) consists of two 'flyer' plates sandwiched around an explosive filling. Upon detonation, typically initiated by a penetrating threat, the explosive drives the flyer plates apart. When oriented off-normal to the impacting threat, the imparted motion moves the flyer plates across the threat, continually feeding new armour material into the interaction, see Fig. 1. Conventionally, ERA is used to reduce the penetration of shaped charge jets (SC) such that the residual jet, following interaction with the ERA, can be arrested by a vehicle's passive armour array [1]. ERA can also be designed primarily for kinetic energy (KE) projectiles, requiring much thicker flyer plates than those used for SC designs [2,3].

The effectiveness of ERA depends on the flyer plate and explosive material, thickness, geometry (i.e., length), orientation, and imparted velocity, as well as characteristics of the threat such as velocity, diameter, length, etc. For fielded system, additional system engineering variables also influence performance, e.g., cassette design, etc. The key mechanisms that influence the performance of ERA systems are well described in a series of publications by Held (e.g. Ref. [4], etc.), Mayesless (e.g. Ref. [5], etc.) and co-authors. Due to the difficulty and cost associated with performing and adequately instrumenting explosive trials, as well as security restrictions, there have been limited public studies that report on specific ERA experiments and their performance. Mayesless et al. [6] reported on experiments designed such that only the front or rear flyer plate interacted with the threat (a shaped charge jet) to understand the main phenomena involved. Brown and Finch [7] performed experiments with bare explosive sheets and ERA, impacted by shaped charge jets. Ismail et al. [8] performed a parametric evaluation of ERA orientation, flyer plate and explosive layer thickness, and point of impact. Lidén [9] investigated the performance of ERA against long rod penetrators and shaped charge warhead in reverse ballistic tests and explosive firings. Lanz et al. [10] examined the influence of explosive type on ERA

* Corresponding author.
E-mail address: shannon.ryan@deakin.edu.au (S. Ryan).
Peer review under responsibility of China Ordnance Society

https://doi.org/10.1016/j.dt.2024.05.007
2214-9147/

Please cite this article as: P. Moldtmann, J. Berk, S. Ryan *et al.*, Adaptive optimisation of explosive reactive armour for protection against kinetic energy and shaped charge threats, Defence Technology, https://doi.org/10.1016/j.dt.2024.05.007

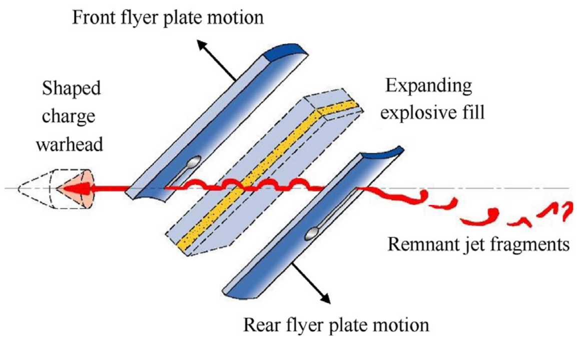

Fig. 1. Disruption of a shaped charge jet by explosive reactive armour (ERA) after [3].

performance and collateral damage. Multiple studies have investigated the performance of non-metallic flyer plates to reduce potential collateral damage, e.g. Refs. [11–13], etc.

Numerical methodologies have been well utilised for further understanding the key mechanisms of ERA performance, e.g., El-Shenawy et al. [14] numerically evaluated the effect of flyer plate and explosive layer thickness on ERA performance at a fixed angle and Li et al. [15] investigated the performance of flat and V-shaped ERA with experimentally-validated simulations. Becker et al. [16] performed many ERA simulations, the results of which were used to train an artificial neural network (ANN) that could predict the position and deformation of steel flyer plates. Moldtmann et al. [17] developed numerical models of ERA with high hardness armour steel and composite flyer plates interacting with long rod KE projectiles and shaped charge jets and evaluated their performance against experimental measurements.

To our knowledge systematic optimisation studies for this armour technology have not been performed in the open literature, although Held performed a trial-and-error optimisation of shaped charge design for robust performance against ERA in [18] and Mickovic et al. developed an analytical model for the expressed purpose of optimising an ERA design in [19]. In this study we utilise Bayesian optimisation (BO), a statistical method that can be used to efficiently optimise expensive black-box functions, to design an ERA for protection against a kinetic energy projectile and shaped charge jet. Here 'expensive' means that data is costly to generate, either in terms of time or other resources, and 'black box' means that the function to be optimised is unknown and only allows point-wise evaluations (i.e., experimental sampling). The utility of BO has been demonstrated across a wide range of applications, traditionally in machine learning [20] but more recently in applied science and engineering domains such as robotics [21], material discovery [22], and material constitutive modelling [23]. In Ref. [24] BO was used as part of a novel human-BO teaming paradigm to optimise the design of a multi-layer armour for protection against two small calibre KE threats. The objective of this study is to compare the optimisation performance of conventional BO with the human-BO teaming paradigm introduced in [24] and a trial-and-error methodology employed by a human domain expert, applied to the design of an ERA.

2. Problem definition

The objective of the optimisation study is to design a minimum weight ERA that reduces the residual penetration of a surrogate medium calibre KE long rod penetrator and a surrogate PG-7 SC warhead into a rolled homogeneous armour steel (RHA) witness block to less than 40 mm. The optimisation activity will use numerical simulation data, see Section 3 for details. The KE penetrator used in the simulations is based on a laboratory round from the French-German Research Institute of Saint Louis (ISL), shown in Fig. 2, with a nominal impact velocity of 1400 m/s and a reference penetration into a semi-infinite RHA steel block of 98 mm. The SC used in the simulations is based on a unitary 79 mm calibre charge from Dynamit Nobel Defence (DND Experimentalhohlladung Typ B), a schematic of which is provided in Fig. 3, with a reference penetration into a semi-infinite RHA steel block of 350–380 mm.

The ERA consists of a front (1st) flyer plate, explosive layer, and rear (2nd) flyer plate, a schematic of which is provided in Fig. 4. A witness block of MARS380, located 200 mm behind the rear surface of the ERA, measured along the projectile line-of-flight, is used to record the residual depth of penetration. Material options and respective thicknesses for each layer of the ERA are summarised in Table 1. Within the defined design matrix there are a total of 1792 potential ERA configurations.

3. Numerical modelling

3.1. Setup of the numerical model

The ERA designs were numerically evaluated against the KE and SC threats using the IMPETUS Afea solver ®, a non-linear explicit finite element code. Specifically, we used the γ-SPH add-on from ABSTRAO, a novel meshless method based on Arbitrary Lagrangian Eulerian (ALE) considerations (see Ref. [25] for details). The ERA plates and witness block are modelled with 200 mm vertical extension and 100 mm lateral extension (although half symmetry was used so only a 50 mm wide block was simulated). ERA plates at obliquity were sized as required to provide the 200 mm vertical coverage. The witness block was modelled as an 80 mm thick MARS380 plate. Representative still images of the simulation models are provided in Fig. 5 for the KE and SC threats.

The jet formation is not explicitly modelled in each of the ERA vs. SC simulations. Rather, to save computational time, we have run the jet formation as a prior model and import it into the ERA simulations. The initial SC model is run for 40 ms, during which the jet tip propagates approximately 170 mm (equivalent to the ideal standoff of the 79 mm calibre SC warhead used in the experiments). The explosive was detonated in the simulations based on a time delay corresponding to when the KE or SC penetrators were estimated to first encounter the explosive layer after penetrating the front flyer plate. The time delay was calculated based on the geometry of the threat, impact velocity, and line-of-sight thickness of the ERA front flyer plate.

To complete the project in the allotted time it was required that individual simulation run times were on average less than 90 min, assuming full utilisation of all available IMPETUS licenses (i.e., 4 runs in parallel, each using a full GPU and associated VRAM). This time requirement put a hard limit on the resolution that could be used in the simulations. Balancing the competing objectives of accuracy and run time, we used SPH elements with 0.5 mm spacing for the KE and SC threats and elements with 0.8 mm for the ERA plates and witness block. Details on the mesh resolution vs. accuracy trade-off study are provided in Ref. [26], together with additional validation and trade-off studies, e.g., attempts to reproduce reference penetration of the two threats, etc.

A summary of the material models and model parameters used in the simulations is provided in Table 2. Specifically, we use the following for all simulations.

- WHA KE penetrator: Mie-Gruneisen equation of state (EoS) with parameters from Ref. [27], Zerilli-Armstrong strength model with parameters from Ref. [28], and Cockroft-Latham failure model with parameters from the IMPETUS material library.

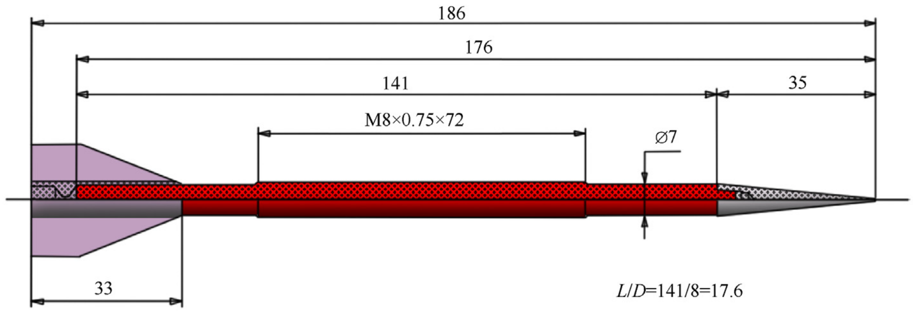

Fig. 2. Details of the surrogate KE penetrator [17]. Units in mm.

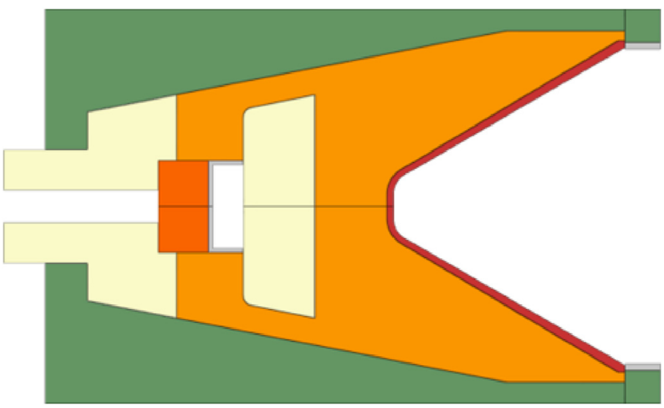

Fig. 3. Schematic of the surrogate SC warhead. The copper liner is shown in red, main explosive charge in orange, and booster in dark orange.

- Copper SC jet: Mie-Gruneisen EoS with parameters from Ref. [29] and a modified Johnson-Cook strength model (after [30]) with parameters from the IMPETUS material library. No failure model is used.
- ARMOX500 flyer plates: linear EoS, Johnson-Cook strength model with parameters from Ref. [31] and Cockroft-Latham failure model with parameters from the IMPETUS material library.
- Dyneema HB26 flyer plates: as anisotropic fabric models are not yet implemented in IMPETUS for SPH-resolved parts, the Dyneema plates were modelled as a stack of unjoined, 2 mm thick isotropic plates to approximate the low through thickness strength of Dyneema laminates (relative to their in-plane strength), after the sub-laminate methodology defined in Ref. [32]. These isotropic sublaminates are modelled using a linear EoS, Johnson-Cook strength and the Cockroft-Latham failure model with parameters calibrated to agree with the

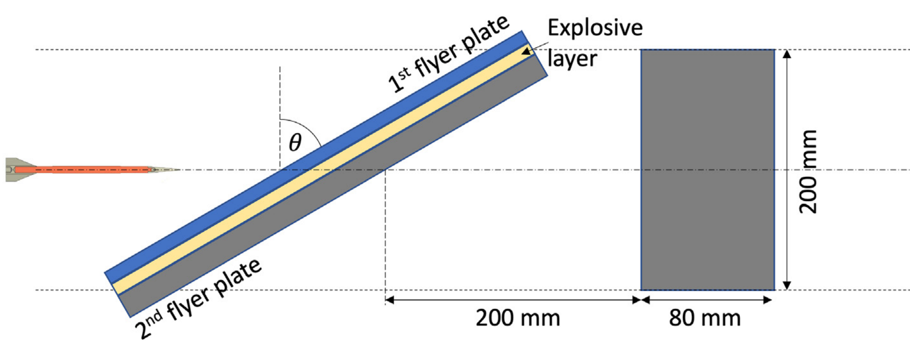

Fig. 4. Schematic of the ERA target, with front (1st) and rear (2nd) flyer plates sandwiched around an explosive layer. A 200 mm air gap is measured between the rear surface of the 2nd flyer plate and the front surface of the witness block. The ERA orientation is measured between the surface of the front flyer plate and normal to the threat line-of-flight.

Table 1
Details of the design matrix.

Feature	Material	Configurations	No. of options
1st flyer plate	ARMOX500	4–10 mm thick in 2 mm increments	4
	Dyneema HB26	10–40 mm thick in 10 mm increments	4
Explosive layer	Semtex	2–8 mm thick in 2 mm increments	4
2nd flyer plate	ARMOX500	4–10 mm thick in 2 mm increments	4
	Dyneema HB26	10–40 mm thick in 10 mm increments	4
Orientation	n/a	0–60° in 10° increments	7
Air gap[a]	n/a	200 mm [fixed]	1
Witness block	MARS380	80 mm [fixed]	1
		Total:	**1792**

[a] measured from rear surface of 2nd flyer plate to front surface of witness plate along the shot axis.

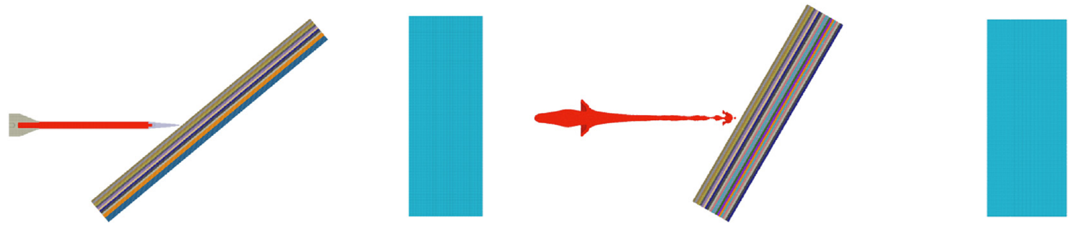

Fig. 5. Visualisation of the simulation models for ERA with Dyneema flyer plates impacted by a KE long rod penetrator (left) and shaped charge jet (right). All parts are discretised using the γ-SPH plug-in in IMPETUS.

Table 2
Constitutive models and model constants used in the numerical simulations.

Physical and thermal properties

Material	ρ/(kg·m^{-3})	E/(MPa)	ν/(−)	χ/(−)	C_p/(J·kg^{-1}·K^{-1})
ARMOX500	7850.0	207.0×10^3	0.30	0.9	455.0
Dyneema HB26	980.0	1.0×10^3	0.45	0.9	
MARS380 [33]	7840.0	210.0×10^3	0.29	0.9	452.0
WHA [27]	17700.0	340.0×10^3	0.28	0.9	155.0
Copper	8900.0	115.0×10^3	0.3	0.9	

Gruneisen Equation of State (EoS)

Material	S_1/(−)	Γ/(−)
WHA [27]	1.237	1.54
Copper [29]	1.489	1.99

JWL Equation of State (EoS)

Material	A/(GPa)	B/(GPa)	R_1/(−)	R_2/(−)	ω/(−)	E_0/(kJ·cm^{-3})
Semtex [34]	759.9	12.56	5.1	1.5	0.29	7.05

Johnson-Cook strength model

Material	A/(MPa)	B/(MPa)	n/(−)	C/(−)	m/(−)	$\dot{\epsilon}_0$ /(s^{-1})	T_0/(K)	T_m/(K)
ARMOX500 [31]	849.0	1340.0	0.0923	0.00541	0.87	1.0	293	1800
Dyneema HB26	50.0							
Copper[a]	90.0	292.0	0.31	0.3	0.45	1.0	293	1356

Zerilli-Armstrong strength model

Material	σ_g/(MPa)	k_h /(MPa·mm$^{1/2}$)	l/ (mm)	K/ (MPa)	n/ (−)	B/ (MPa)	β_0 /(K^{-1})	β_1 /(K^{-1})	$\dot{\epsilon}_0$ /(s^{-1})
MARS380 [33]	938.3	0.0	1.0	416.7	0.28	2294.5	0.0079	0.0002	1.0
WHA [28]	100.0	0.0	1.0	1047.0	0.09	1381.0	0.0025	0.00011	1.0

Cockroft-Latham failure

Material	W_c /(MPa)	σ_s /(MPa)	t_s /(s)	α_s /(−)	β_s /(−)
ARMOX500	2061.0				
Dyneema HB26	100.0				
MARS380 [33]	2200.0				
WHA	450.0	1.1×10^3	2.25×10^{-6}	1.0	1.0

[a] modified formulation of Johnson-Cook strength model after [30], implemented as MAT_METAL in IMPETUS.

MAT_FABRIC model for Dyneema HB26 in the IMPETUS material library (see section 3.2 for validation of this approach).

- MARS380 witness plate: linear EoS, Zerilli-Armstrong strength model and Cockroft-Latham failure model with parameters from Ref. [33]. Parameters from Ref. [33] were derived for MARS190, an earlier product name for MARS380.
- Semtex explosive layer: JWL EoS with parameters for m/46 from Ref. [34]. m/46 is a Swedish explosive with a similar PETN/RDX content to Semtex PI SE M (HP) [35]. The detonation velocity of the explosive was set to 6684 m/s.

3.2. Validation of the numerical model

The simulation models were validated via comparison with experiments performed at the ISL with the two threats described in section 2. Three ERA configurations were experimentally evaluated, schematics of which are provided in Fig. 6. Details of the experimental setup and instrumentation are provided in Ref. [17]. For the KE experiments the explosive is too insensitive to be initiated by the impact of the projectile. As such, a detonator was installed at the impact area in these tests, triggered based on an estimated time when the penetrator tip reaches the explosive (i.e., consistent with the timing used in the simulations). The experiments used 200 mm wide flyer plates and witness blocks, compared with the 100 mm width used in the simulations.

First, we consider the justification for modelling the Dyneema as a simplified stack of isotropic layers in the γ-SPH discretisation. The motivation for using this discretisation scheme was to reduce the runtime of the simulations. For a representative KE simulation, a Lagrangian-discretised projectile and ERA target required approximately 14 h to run to completion (380 μs) on an Nvidia RTX A5000

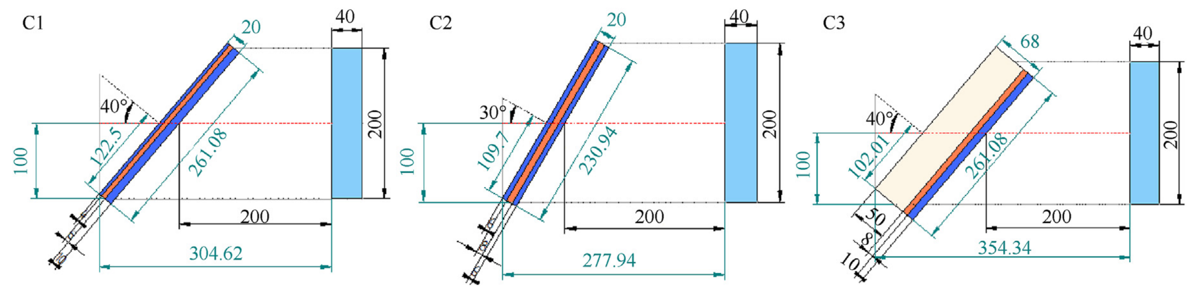

Fig. 6. ERA configurations experimentally evaluated for validation of the numerical simulations. Impact is from left-to-right. Configuration C1: 4 mm ARMOX500, 6 mm Semtex, 10 mm ARMOX500 at 40° NATO angle, C2: 6 mm ARMOX500, 8 mm Semtex, 6 mm ARMOX500 at 30° NATO angle, C1: 50 mm Dyneema HB26, 8 mm Semtex, 10 mm ARMOX500 at 40° NATO angle. All experiments use a stack of 5×40 mm thick MARS380 witness blocks. Units in mm.

GPU. The same model, using a γ-SPH discretisation for the KE projectile and ERA target, required approximately 2.5 h to run on the same computing hardware. Simulations using the SPH and Lagrangian representations of the Dyneema laminate are compared with an experimental radiograph in Fig. 7 (the ERA configuration is C3 from Fig. 6). We can observe that the simulation using a γ-SPH discretisation shows incipient fracture of the rod towards the tail, while the simulation using a Lagrangian discretisation shows incipient fracture closer to the nose (and then perhaps again towards the tail). At 120 μs post-impact the experimental image does not show any clear fracture, although bending is apparent. By 160 μs post-impact the experimental image shows fracture of the rod at the tail (see Ref. [17]). Based on this comparison we find that both discretisation schemes result in reasonable qualitative agreement with the experiment, and that the Lagrangian model with a more representative anisotropic material formulation does not provide any improvement in agreement with experiment.

Using the γ-SPH discretisation, we compare the residual penetration measurements and numerical predictions for the three ERA configurations from Fig. 6 in Table 3. We can observe that KE residual penetration was significantly under predicted in simulations of ERA configurations with ARMOX500 for the front and rear flyer plates (C1 and C2), and reasonably well predicted in the simulation of configuration C3 which had a Dyneema HB26 front flyer plate and an ARMOX500 rear flyer plate. The SC residual penetration was also significantly under predicted in the simulations of configurations with front and rear ARMOX500 flyer plates (C1 and C2), and reasonably well predicted for configuration C3 with a Dyneema HB26 front flyer plate and an ARMOX500 rear flyer plate. For the shaped charge threat against the C1 and C2 ERA simulations the 80 mm witness block was completely perforated by the residual jet. In hindsight a thicker witness block to capture the full residual penetration of the SC threat would have been preferrable, however, this was impacted by the run-time vs. numerical model accuracy trade-off previously discussed. Although the simulations are not quantitatively accurate in predicting the residual penetration for either the KE or SC threat, they do capture the respective ranking of the three evaluated configurations. In Fig. 8 we plot the residual penetration measurements and simulation predictions for the three reference ERA configurations against the KE and SC threats. The performance ranking for the three configurations against the KE threat is determined in the experiment as C1 > C2 > C3 and against the SC threat is determined as C3 > C2 > C1 (i.e., the inverse of that for the KE threat). This result is not unexpected as one of the main challenges in the design of armour systems is the different defeat mechanisms for KE and SC threats. In the simulations we see that these rankings are maintained, i.e., C1 > C2 > C3 against the KE threat and C3 > C2/C1 against the SC threat.

In Fig. 9 some representative radiographs from the experiments are compared to equivalent simulation frames for qualitative evaluation. In general, we find that the simulations show quite good agreement with the experiments. The position and shape of the ERA plates is reasonably well predicted, including the Dyneema plates. The erosion, deformation, and fragmentation of the KE rod is reasonably well predicted, although the simulation shows severe rod bending that is not representative of the experimental image. The tip velocity of the SC model appears to be significantly higher in the experiments than the simulation, while the simulation seems to overpredict the jet diameter. The jet is also shown to be much too coherent in the numerical simulations through the interaction with the ERA, compared to the experiment.

The models are considered to provide reasonable qualitative accuracy for use in the optimisation study. Improvements in the predictive accuracy are likely achievable with increased resolution of the SPH elements while others require modifications to the material constitutive models/model constants (e.g., WHA fracture criteria, etc.) and solver (e.g., treatment of SPH elements lining a penetration cavity, jet viscosity, etc.). Increases in SPH resolution,

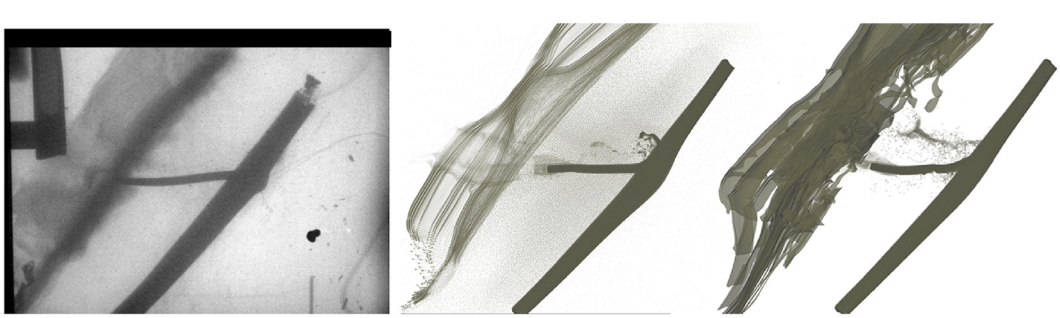

Fig. 7. Comparison between ERA with 10 mm thick Dyneema front plate and 10 mm thick ARMOX500 rear plate impacted by the KE rod at 120 μs post-impact. Left: Experimental radiograph, middle: γ-SPH discretisation and, right: Lagrangian discretisation for the Dyneema laminate.

Table 3
Comparison of the experimental measurements and numerical predictions for the 3 ERA configurations tested at the ISL (C1, C2, C3 from Fig. 6). For the simulations an 80 mm thick MARS380 witness block was used, thus residual penetrations >80 mm could not be measured.

Threat	ERA design					Residual DoP /(mm)			Error /(%)
	#	Front plate	Explosive	Rear plate	Orientation	Exp	Sim	Error	
KE	C1	4 mm ARMOX500	6 mm Semtex	10 mm ARMOX500	40°	21	2	−19	89
SC	C1	4 mm ARMOX500	6 mm Semtex	10 mm ARMOX500	40°	119	>80[a]	−39	33[a]
KE	C2	6 mm ARMOX500	8 mm Semtex	6 mm ARMOX500	30°	25	18	−7	29
SC	C2	6 mm ARMOX500	8 mm Semtex	6 mm ARMOX500	30°	101	>80[a]	−21	21[a]
KE	C3	50 mm Dyneema HB26	8 mm Semtex	10 mm ARMOX500	40°	33	39	6	18
SC	C3	50 mm Dyneema HB26	8 mm Semtex	10 mm ARMOX500	40°	77	74	−2	3

[a] The simulations had an 80 mm thick witness block that was completely perforated.

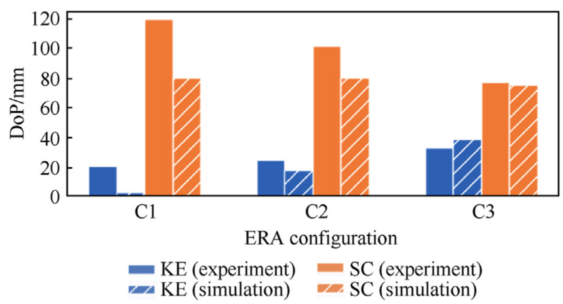

Fig. 8. Comparing the measured and simulated residual depth of penetration for the three ERA configurations subject to experimental firings. Note that the simulations of CI and C2 with the SC threat resulting in complete perforation of the 80 mm MARS380 witness plate (i.e., the actual simulated penetration would have been greater than 80 mm in these two models).

however, come with increased computational cost, which was not feasible for this study.

3.3. *Application of the numerical model*

The numerical model was applied to simulate the entire design space defined in Table 1 against both the KE and SC threats prior to performing the optimisation. This was done to facilitate the evaluation of the different optimisation methodologies and avoid 'contamination' of the human experts (e.g., providing them access to some of the numerical simulation results and then making them re-start the optimisation due to problems encountered with the code). Simulations were performed in parallel on computing facilities at the ISL and Deakin-A2I2, using 3 different GPUs: NVIDIA RTX A5000, NVIDIA RTX A6000, and NVIDIA DGX A100.

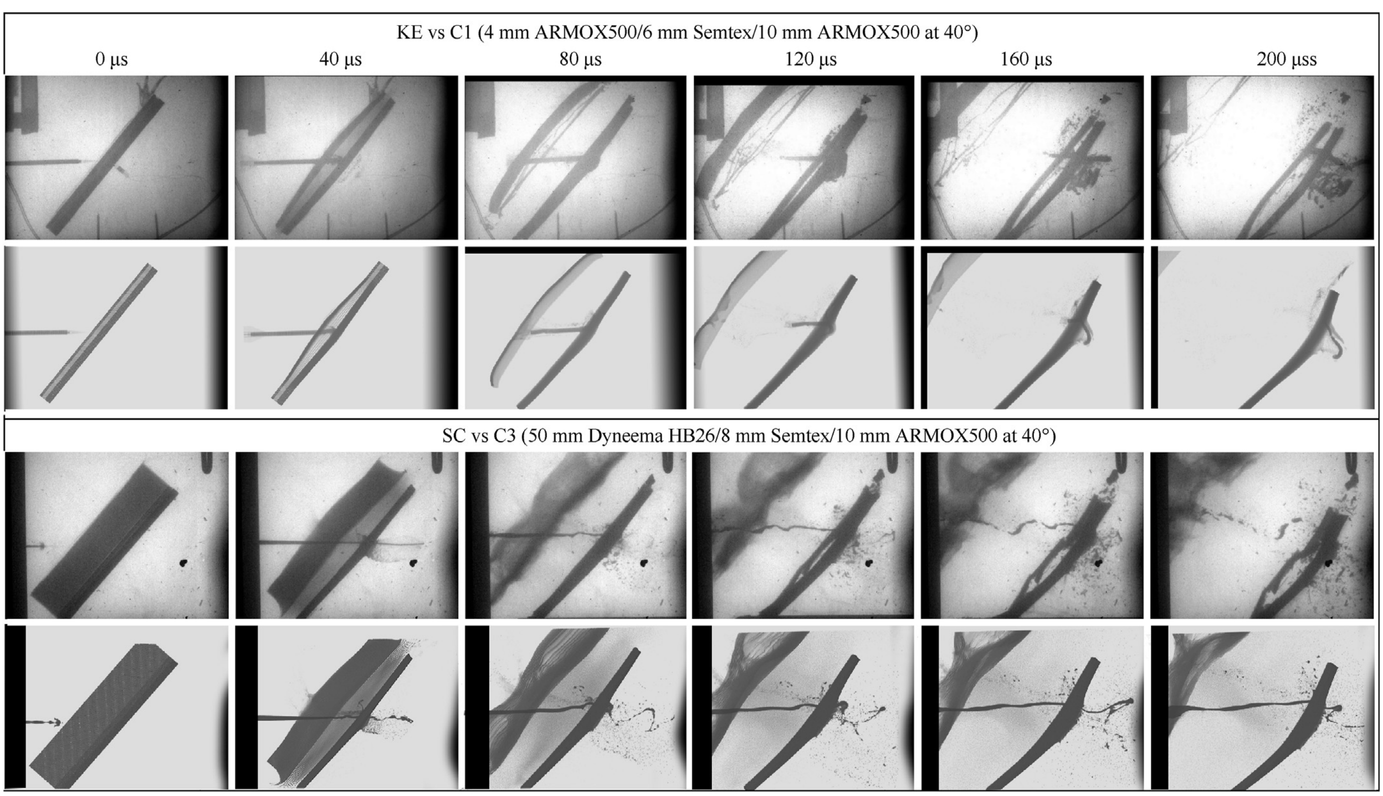

Fig. 9. A comparison between experimental radiographs (upper) with simulation frames (lower) in 40 µs intervals for impact of the kinetic energy surrogate threat (top 2 rows) and shaped charge surrogate threat (bottom 2 rows) against ERA configuration C1 (top 2 rows) and C3 (bottom 2 rows).

4. Design optimisation strategies

We perform the design optimisation using three different strategies: (1) a conventional trial-and-error methodology employed by a terminal ballistics expert at the ISL, (2) a machine learning-based adaptive sampling methodology, Bayesian optimistion (BO), and (3) a novel human-machine teaming framework that partners an engineer from the ISL (different from the expert used for approach #1) with a Bayesian optimiser. All three optimisations use the same dataset but are performed independently of one another, e.g., the two humans are not allowed to communicate with each other during the study, or review results of the other approaches.

4.1. Trial-and-error ("Expert")

Traditional armour design is performed by engineers who utilise their experience and domain knowledge to formulate an experimental design, test it, and iterate upon it based on the experimental result. Experts may utilise additional sources of information such as reference publications, numerical solvers, or semi-analytical equations, to improve their knowledge in specific design regimes. This knowledge is organised in complex schema containing concepts, attributes, and relationships [36].

The role of the human expert was fulfilled by an experienced researcher from the ISL with a diploma in civil engineering and a PhD in experimental mechanics. At the time of this study, they had 13 years of terminal ballistics experience, including approximately four years of experience working with ERA (primarily for protection against kinetic energy projectiles). The expert has experience with performing numerical simulations and ballistic experiments with both kinetic energy projectiles and shaped charge warheads.

4.2. Bayesian optimisation ("Autonomous BO")

BO has been demonstrated to find a system optimum in fewer experiments than other black box optimisation techniques such as e.g., evolutionary algorithms, Tree Parzen estimator, etc. [37,38]. BO uses a model-based approach with an adaptive sampling strategy to minimise the number of experiments required. The black box function is typically modelled with a Gaussian Process (GP), a regression model that simultaneously provides the predicted mean, $\mu_t(x)$, and the epistemic uncertainty, $\sigma_t(x)$ at any point x in the input space, given a set of observations $\mathscr{D}_{1:t} = \{(x_1, y_1), (x_2, y_2), \ldots (x_t, y_t)\}$ where x_i is a combination of design input variables and y_i is the corresponding system output. Adaptive sampling [39] is a process that progressively determines the ideal conditions for experimentation by balancing two criteria: *exploration* − to reduce the uncertainty of the surrogate function of the system, and *exploitation* − to maximise the value of the optimisation objective. These objectives are analogous to the evaluation of designs near known, good solutions (exploitation) and the evaluation of novel designs that are unlike those that have previously been evaluated (exploration). A schematic of the Bayesian optimisation workflow is provided in Fig. 10.

For this project we utilised a constrained, single-objective BO methodology. A GP surrogate model was used to predict the probability that the ERA designs would meet the DoP constraint of <40 mm for each of the two threats. The objective was to minimise the line-of-sight weight of the ERA, represented by the areal density (AD), for which an analytical solution exists, i.e., given an ERA design, the AD can be analytically computed. The acquisition function was therefore a product of the expected improvement in terms of AD reduction over the best-known solution and the probability that the design would meet the probabilistic constraint. This differs from conventional BO in which both the objective and constraint would both be modelled by a GP. To initialise the optimisation, we randomly selected three designs from the grid.

4.3. Human-machine teaming ("BO-Muse")

Standard BO assumes no initial knowledge of the system, i.e., the GP surrogate model is described by a Gaussian prior of $\mu_t(x) = 0$ and $\sigma_t(x) = 1$ across the entire design space. The GP is initialized by a limited number of experiments/sample points, in this study randomly sampled from the design space. The optimisation progresses through the iterative loop shown in Fig. 10 typically with no outside input. It is hypothesised that a human with domain knowledge could accelerate this process. However, complexities involved in extracting knowledge from the human in a form that can be exploited by the BO mean that this remains largely unproven. Some efforts have been made to incorporate domain-specific knowledge in BO [41] or transfer learning from previous experiments [42] with limited success.

In [43] a human-BO collaborative framework is proposed in which a human's suggestions are incorporated in parallel with that of the BO acquisition function during the adaptive sampling step in Fig. 10. In the balancing of optimisation objectives introduced in section 4.2 humans have been shown to be highly skewed towards exploitation [44]. When the task requires specialised human experts, as might be expected for an applied engineering problem like armour design, the balance swings further towards exploitation. External stimuli have been shown to boost human creativity, possibly providing a means to escape this exploitation-skewed dilemma [45] known as cognitive entrenchment. BO-Muse is a formal framework for inserting BO into a human expert's workflow, i.e., the optimisation is 'led' by the human and the BO-generated suggestions are intended to stimulate their creativity, breaking them free from their entrenched, exploitation-heavy, sampling strategy. The BO-Muse workflow is schematically shown in Fig. 11.

The objective function used in the BO-Muse methodology was identical to that used for the autonomous BO methodology, described in section 4.2. In terms of the BO-Muse framework this means that the BO suggestions were based on a 'balanced' objective of maximising the improvement in the best-found ERA solution (i.e., maximise the reduction in weight) and providing the largest improvement in the surrogate model (i.e., minimise the uncertainty in the GP surrogate model). This is different to the strategy used in Ref. [24], for example, which used an 'exploratory' strategy for the BO. The BO-Muse optimisation was performed in batches of two: one suggestion from the human and one suggestion from the BO.

The role of the human in the BO-Muse team was fulfilled by a doctoral student from the ISL with a Master's degree in engineering. At the time of this study, they had 3 years of terminal ballistics experience, albeit none working directly with ERA. They have experience performing numerical simulations and ballistic experiments with kinetic energy projectiles and fragments and a basic knowledge of the general working principles of ERA.

5. Results

As the entire design grid was simulated prior to the optimisation we can definitively list the best valid solutions according to the numerical model (where valid indicates that the ERA met the residual penetration requriement for both threats). The top 5 valid solutions are provided in Table 4 together with the best flyer plate material-specific valid solutions.

Given the quantitative inaccuracy of the simulation models (see section 3.2) it is difficult to state whether the results shown in Table 4 are true. However, we can assess their reasonableness by considering the experimental findings reported in section 3.2 and

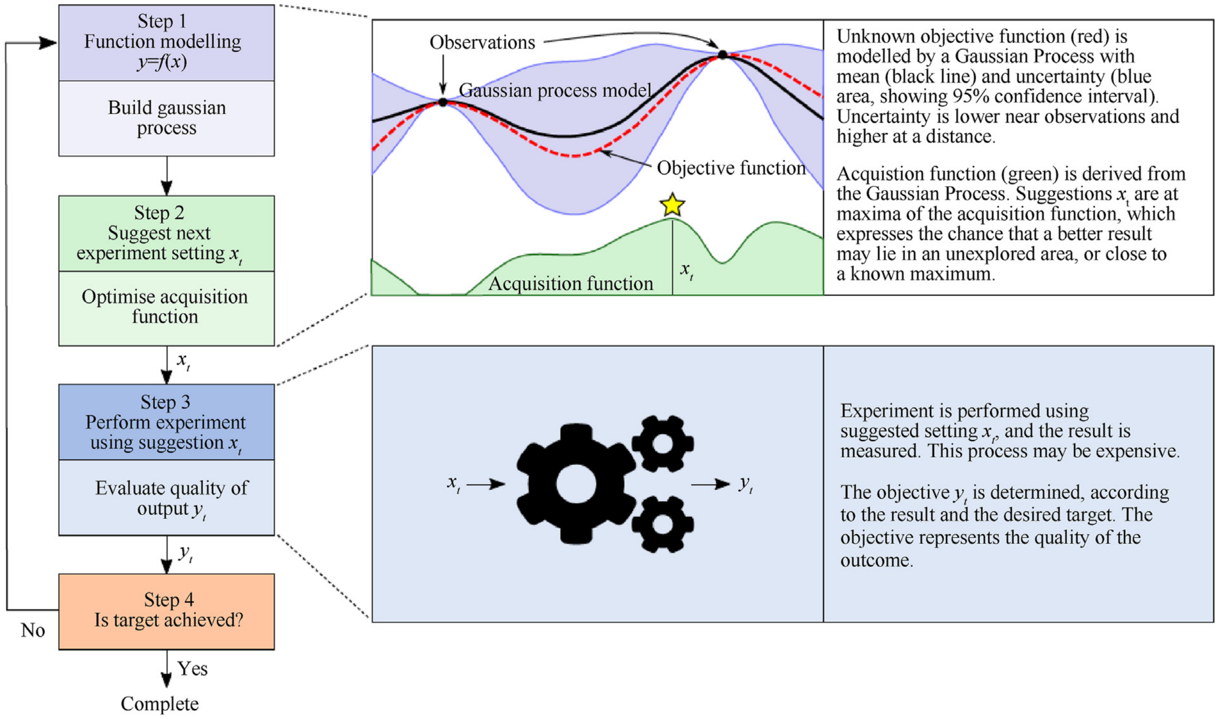

Fig. 10. An overview of Bayesian optimisation from Ref. [40]. An unknown function is approximated by a Gaussian process (GP), from which an acquisition function makes suggestions for the next experimental condition to be evaluated, x_t. The experimental results, y_t, are used to update the GP surrogate model, and an iterative, adaptive sampling loop is executed until the optimisation objective is achieved or the experimental budget is exhausted.

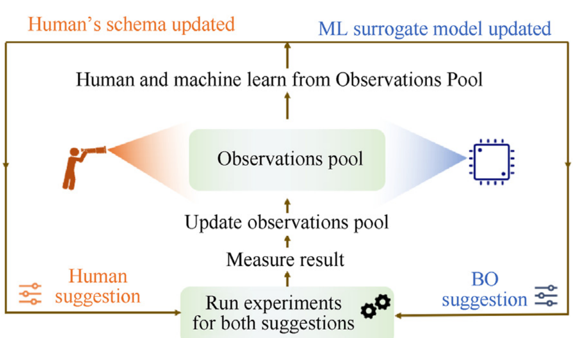

Fig. 11. A framework for collaborative teaming of humans with domain knowledge and Bayesian optimisation after Ref. [43], referred to as BO-Muse, in which adaptive sampling is performed according to suggestions by both the human and BO acquisition function.

Ref. [17]. From Fig. 8 we find that the C3 configuration, i.e., the only experimental configuration with a Dyneema HB26 flyer plate, provided the best residual penetration results of the three configurations tested. The line-of-sight AD of the three experimental configurations was 154, 121, and 168 kg/m^2 for the C1, C2, and C3 configurations, respectively. Thus, the best performing configuration, C3, was also the heaviest. Nonetheless, for a modest increase in areal weight between the C1 and C3 configurations (14 kg/m^2 or 8% increase) a significant reduction in the SC residual DoP was achieved (40 mm or 35%). For the majority of the simulations, we observed that the SC was the dominant threat (i.e., higher residual penetration than the KE threat). The effectiveness of the Dyneema HB26 against the SC observed in the C3 experiments, therefore, suggests that the optimum results identified from the simulations and reported in Table 4 are reasonable.

Table 4
Details of the optimum ERA designs (i.e., minimum line-of-sight areal density with residual penetration for both threats <40 mm) according to the numerical simulations.

Rank	Angle/ (°)	Layer 1		Explosive	Layer 3		AD_{LoS} /(kg·m^{-2})	Result, DoP_{WP} /(mm)	
		Mat.	t /(mm)	t /(mm)	Mat.	t /(mm)		KE	SC
1	50	Dyneema HB26	10	8	Dyneema HB26	10	47.9	38.5	34.7
2	60	Dyneema HB26	10	6	Dyneema HB26	10	56.0	37.9	34.9
3	60	Dyneema HB26	10	8	Dyneema HB26	10	61.6	36.5	26.8
4	50	Dyneema HB26	10	8	Dyneema HB26	20	63.2	34.8	34.2
5	40	Dyneema HB26	10	8	Dyneema HB26	30	65.8	34.0	37.2
10	50	Dyneema HB26	10	8	ARMOX500	4	81.5	20.5	34.9
18	50	ARMOX500	4	8	Dyneema HB26	20	96.8	31.7	30.3
43	50	ARMOX500	4	8	ARMOX500	4	115.1	28.7	32.5

5.1. Expert

During the optimisation the expert would select an ERA design based upon their experience, incorporating any additional sources of information available such as reference publications, semi-analytical models, etc. The simulation results for that design would be provided to the expert in the form of residual DoP values for each threat, together with still frames and videos from the simulations. The expert would review the result and then select the next design for evaluation.

The expert achieved a best solution weight of 81.5 kg/m^2 in 11 evaluations, the results for which are provided in Table 5. During the optimisation the expert was asked to document the motivation behind their designs, so we have some insight into their thought process. We can observe that, in general, the expert changed one input variable in each iteration, e.g., the thickness of the explosive layer, the front flyer plate material, the flyer plate thickness, etc. Based on their domain knowledge, their design basis was aimed to achieve a relatively heavy rear flyer plate with low velocity for effective performance against the KE projectile and a relatively light front flyer plate with high velocity for for effective performance against the SC jet. Additionally, based on an understanding that ERA performance is typically proportional to impact obliquity, high obliquity design options were preferred.

Although directed to evaluate up to as many as 50 designs, the expert stopped after the 11 shown in Table 5 as this is what they would have done in "a real study". After the conclusion of the study the expert was provided with details of the optimum ERA solution, to which they stated that they would not have considered Dyneema HB26 for both flyer plates as they had never seen such a design in literature or other research projects.

5.2. Autonomous BO

As BO can be sensitive to the initialisation condition, 20 runs of the autonomous BO strategy were performed. Each optimisation ran for 30 iterations. In 16 of the 20 runs the optimal solution with an AD $=$ 47.9 kg/m^2 was identified. The other four runs failed to identify a single valid solution in the 33 designs evaluated (3 initial designs $+$ 30 iterations). The median number of design iterations required to identify the optimum solution was 15 across the 20 optimisation runs. A comparison of the 20 optimisation runs is shown in Fig. 12. The median and mean of the best-found result, calculated at each iteration for the 20 optimisation runs, is also plotted, together with the upper (75%) and lower (25%) quartile range.

5.3. BO-Muse

The BO-Muse strategy found the optimal solution on the 5th iteration (10th sample). The results of the BO-Muse optimisation are given in Table 6. We can observe that the optimal solution was identified by the BO. We also can observe that the human expert, like the expert in section 5.1, typically changes one variable at a time, e.g., the thickness of the ARMOX500 rear flyer plate, the thickness of the front Dyneema HB26 flyer plate, etc. The expert does not appear to have adjusted their strategy much as a result of the BO suggestions, although it is difficult to verify as the expert did not provide the motivation behind their designs during the optimisation activity. We can observe some similar behaviour to the other human expert, e.g., utilising their domain knowledge to provide a design that has high obliquity, a lighter front flyer plate (for high velocity) and a heavier rear flyer plate (for lower velocity).

6. Evaluation

The performance of the different optimisation methodologies is compared in Fig. 13. The BO-Muse methodology is shown to identify the optimum solution in the least number of sample evaluations (10). The autonomous BO methodology requires, per its median performance, an additional 5 sample evaluations, but also fails to identify any valid solutions in 4 of the 20 optimisation runs. The four autonomous BO runs that failed to find the optimum solution can be explained by the large number of design variants which had a simulated DoP >80 mm for the SC threat (1187/1792 variants). In these four runs not one design was identified that met the penetration criteria over the 30 iterations plus 3 random initial designs. If BO is unlucky and no valid points are identified during the initialisation then the surrogate model has no valid points to regress, resulting in poor performance. This is a strength of the BO-Muse framework as experts are good at identifying valid solutions, even heavy ones, that will provide the necessary initial points for the GP surrogate model to function. The expert's trial-and-error methodology is shown to identify lighter weight solutions more quickly than autonomous BO, but is unable to identify the system optimum, instead ending with a best-found solution of 81.52 kg/m^2 (70% heavier than the actual system optimum). By evaluating the expert's documented design motivation, we can identify clear 'cognitive entrenchment', as expected for an expert with considerable domain expertise. The BO-Muse strategy, which shows comparable performance to the human expert for the initial ~7 samples, successfully overcome this limitation.

The objective of the study was to evaluate the performance of

Table 5
Results of the expert/trial-and-error design optimisation. The best-found solution is shown in bold.

#	Angle /(°)	Layer 1		Explosive	Layer 3		AD_{LoS} /($kg \cdot m^{-2}$)	Result, DoP_{WP} /(mm)		Valid result?
		Mat.	*t* /(mm)	*t* /(mm)	Mat.	*t* /(mm)		KE	SC	
1	60	ARMOX500	4	4	ARMOX500	8	199.6	9.9	51.3	N
2	60	ARMOX500	4	6	ARMOX500	8	205.2	23.0	23.4	Y
3	60	Dyneema HB26	20	6	ARMOX500	8	181.6	7.8	17.8	Y
4	60	Dyneema HB26	20	4	ARMOX500	8	176.0	6.0	42.1	N
5	60	Dyneema HB26	10	4	ARMOX500	8	156.4	3.0	31.8	Y
6	60	Dyneema HB26	10	4	ARMOX500	6	125.0	3.9	44.2	N
7	60	Dyneema HB26	10	6	ARMOX500	6	130.6	10.7	29.5	Y
8	50	Dyneema HB26	10	6	ARMOX500	6	101.6	4.6	48.0	N
9	50	Dyneema HB26	10	8	ARMOX500	6	105.9	4.7	36.8	Y
10	**50**	**Dyneema HB26**	**10**	**8**	**ARMOX500**	**4**	**81.5**	**20.5**	**34.9**	**Y**
11	40	Dyneema HB26	10	8	ARMOX500	4	68.4	10.6	56.2	N

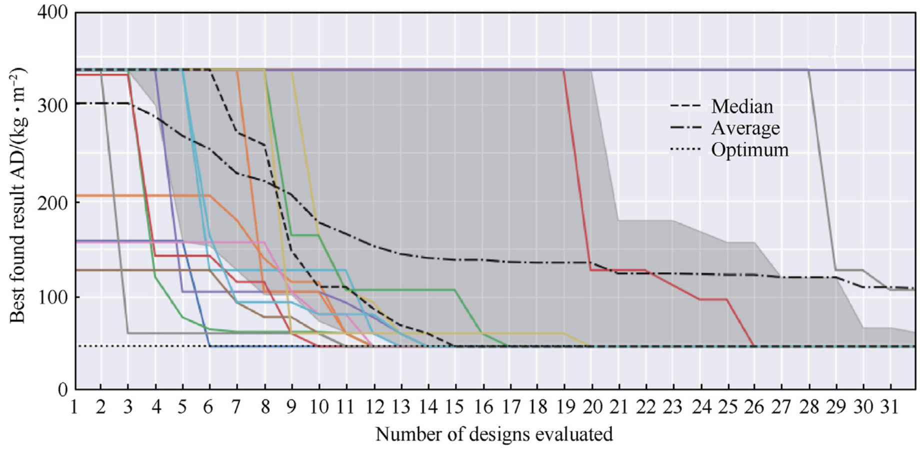

Fig. 12. Results of 20 runs of the autonomous BO methodology. Each optimisation is initialized with three random samples from the design grid and run for 30 iterations. The median result is shown by the long black dashed line, the mean result is shown by the dotted line, and the upper and lower quartile is shown by the shaded region.

Table 6
Results of the BO-Muse design optimisation. The best-found solution is shown in bold font.

#	Source	Angle /(°)	Layer 1		Expl.	Layer 3		AD_{LoS} /(kg·m^{-2})	Result, DoP_{WP} /(mm)		Valid result?
			Mat.	t /(mm)	t /(mm)	Mat.	t /(mm)		KE	SC	
1	Human	60	Dyneema HB26	40	4	ARMOX500	6	183.8	10.6	24.0	Y
2	BO	0	Dyneema HB26	10	2	Dyneema HB26	10	22.4	50.4	80.0	N
3	Human	60	Dyneema HB26	40	4	ARMOX500	4	152.4	22.2	27.8	Y
4	BO	40	Dyneema HB26	30	4	ARMOX500	4	86.7	16.6	80.0	N
5	Human	60	Dyneema HB26	20	4	ARMOX500	4	113.2	21.0	49.8	N
6	BO	0	ARMOX500	4	8	Dyneema HB26	10	52.4	46.5	80.0	N
7	Human	60	Dyneema HB26	20	6	ARMOX500	4	118.8	20.6	32.8	Y
8	BO	60	Dyneema HB26	10	8	Dyneema HB26	10	61.6	36.5	26.8	Y
9	Human	60	Dyneema HB26	10	6	ARMOX500	4	99.2	29.2	28.5	Y
10	**BO**	**50**	**Dyneema HB26**	**10**	**8**	**Dyneema HB26**	**10**	**47.9**	**38.5**	**34.7**	**Y**

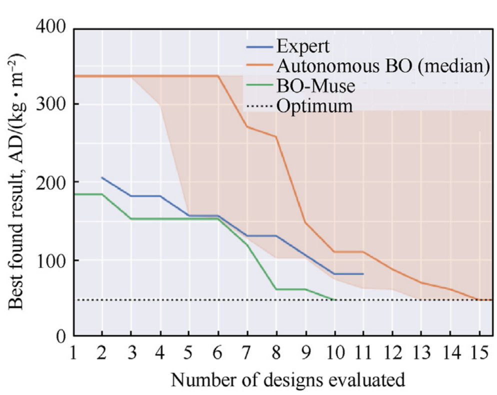

Fig. 13. Comparing the expert, autonomous BO, and BO-Muse optimisation strategies. The expert and BO-Muse optimisations can only be performed once, so the performance curve is deterministic compared to the probabilistic autonomous BO curve, which represents the median (orange line) and upper and lower quartiles (shaded region) from 20 optimisation runs.

different optimisation methodologies on the design problem defined in section 2. For 'simple' experimental design problems in which the number of variables is low it is expected that the trial-and-error methodology employed by human experts is unbeatable due to the ability of humans to conceptualise and simplify complex systems, particularly in the neighbourhood of known solutions. As the complexity of the design space increases, it is hypothesised that methodologies like BO or BO-Muse become more competitive, eventually overtaking the performance of an expert trial-and-error approach. In the design of the threat conditions for this project the goal was to select KE and SC surrogates that provided relatively comparable lethality against ERA designs. Avoiding a dominant threat would mean that the optimisation methodologies would have to consider both threats throughout, rather than down-selecting to a dominant threat, designing a solution to defeat it, and then verifying its performance against the secondary threat. Unfortunately, this was not achieved. Rather, the SC threat was clearly dominant for the vast majority of potential ERA designs in the grid. As such, the optimisation was likely easier than intended, which we expected to favour the expert trial-and-error approach. For an 'easy' optimisation problem in which we expect a human expert to outperform the autonomous BO methodology, we would also expect the BO-Muse methodology to be uncompetitive as the BO suggestions would use up samples that would be less useful than the human equivalents. Comparing the results in Fig. 13 suggests that even though the SC threat was dominant, the design problem was still sufficiently complex for the BO suggestions to positively contribute in the BO-Muse framework.

The numerical simulations were shown in section 3.2 to provide reasonable qualitative agreement with the experiments, including retaining the performance ranking of the three configurations experimentally evaluated against both the KE and SC threats. However, the quantitative accuracy of the simulations was insufficient. As the entire optimisation was performed using simulation data, it is feasible that the inaccuracy of the simulations may disadvantage the human expert if the relative performance of the different designs is not representative of their actual performance. In this event, simulation results may conflict with the existing schema of the expert, built upon more than a decade of experience with experimental measurements. This is a clear limitation of this study. In a more typical optimisation study simulations would be performed on demand, when selected by the human expert or BO algorithm. In such a study there would be orders of magnitude fewer simulations to be performed and, as such, higher resolution simulations with longer run times (and greater accuracy) would be feasible.

7. Conclusions

Three competing optimisation strategies were used to design a minimum weight explosive reactive armour (ERA) to reduce the penetration of a surrogate medium calibre kinetic energy (KE) long rod projectile and a PG-7 surrogate shaped charge (SC) warhead to <40 mm into a semi-infinite witness block of rolled homogeneous armour steel (RHA). The ERA design variables included: material and thickness of the front and rear flyer plates, thickness of the explosive layer, and rotation angle of the arrangement, relative to the flight path of the two threats. The total number of potential ERA designs available for the optimisation was 1792. The optimisation was performed using experimentally validated numerical simulations that were demonstrated to provide qualitatively accurate predictions that maintained relative performance rankings of three experimentally evaluated designs.

A human expert with 13 years' experience in terminal ballistics and armour design applied a trial-and-error optimisation methodology, representative of best practise within ISL and many other large Defence research laboratories. The second optimisation approach utilised an adaptive Bayesian optimisation (BO) methodology with no human input. The third optimisation approach utilised a novel human-BO teaming framework referred to as BO-Muse, utilising a different human with 3 years' experience in terminal ballistics.

The human-machine teaming framework (BO-Muse) was shown to outperform both the stand-alone human and stand-alone machine (BO). Incorporating the human suggestions in the BO-Muse framework was found to improve the convergence rate of the optimisation compared to standard BO, based on the median result across 20 BO optimisation runs. The BO-Muse teaming approach was also demonstrated to improve upon the performance of the human expert employing a traditional trial-and-error methodology, which showed an improved rate of convergence initially compared to standard BO (i.e., found lighter weight solutions faster), but failed to find the true system optima — with a best-found solution 70% heavier than that identified by the BO and BO-Muse methodologies.

Declaration of competing interest

The author (Shannon Ryan) is an Editorial Board Member/Editor-in-Chief/Associate Editor/Guest Editor for *Defence Technology* and was not involved in the editorial review or the decision to publish this article.

References

[1] Held M. Armour. In: Proceedings of the 14th international Symposium on ballistics; 1994. Quebec City.
[2] Held M, Mayseless M, Rototaev E. Explosive reactive armour. In: Proceedings of the 18th international Symposium on ballistics; 1988. Midrand.
[3] Held M. Lectures on armors and warhead mechanisms. In: 23rd international symposium on ballistics, tarragona; 2007.
[4] Held M. Defeating mechanisms of explosive reactive armor sandwiches. In: Proceedings of the 22nd international Symposium on ballistics. Vancouver; 2005.
[5] Mayesless M. Effectiveness of explosive reactive armor. J Appl Mech 2011;78:051006.
[6] Mayesless M, Erlich Y, Falcovitz Y, Rosenberg G. Interaction of shaped-charge jets with reactive armor. In: Proceedings of the 8th international Symposium on ballistics; 1984. Orlando.
[7] Brown J, Finch D. The shaped charge jet attack of confined and unconfined sheet explosive at normal incidence. In: Proceedings of the 12th international Symposium on ballistics; 1990.
[8] Ismail MM, Rayad AM, Alwany H, Alshenawy TA. Optimization of performance of explosive reactive armors. In: Proceedings of the 21st international Symposium on ballistics; 2004. Adelaide.
[9] Lidén E. Interaction between high-velocity penetrators and moving armor components. Uppsala University; 2010. PhD thesis.
[10] Lanz W, Bircher HR, Wyssen D, Bieri G, Rolli MK. Evaluation of high explosive parameters for reactive armour. In: Proceedings of the 19th international Symposium on ballistics. Interlaken; 2001.
[11] Kaufmann H, Koch A. Terminal ballistic effects of low density materials used as confinement plates for explosive reactive armour. In: Proceedings of the 22nd international symposium on ballistics. Vancouver; 2005.
[12] Lidén E, Helte A, Eriksson T. Reactive panels with non-metallic cover plates. In: 30th international symposium on ballistics. Long Beach; 2017.
[13] Hazell P, Lawrence T, Stennett C. The defeat of shaped charge jets by explosively driven ceramic and glass plates. Int J Appl Ceram Technol 2012;9(2):382—92.
[14] El-Shenawy TA, Riad AM, Ismail MM. Penetration and initiation of explosive reactive armors by shaped charge jet. In: Proceedings of the 26th international annual Conference of ICT; 2005. Karlsruhe.
[15] Li X, Yang Y, Lv S. A numerical study on the disturbance of explosive reactive armors to jet penetration. Defence Technology 2014;10(1):66—75.
[16] Becker M, Klavzar A, Wolf T, Renck M. Data-driven prediction of plate velocities and plate deformation of explosive reactive armor. Defence Technology 2022;18(12):2141—9.
[17] Moldtmann P, Limido J, Klavzar A, Ryan S. ERA targets against different threats: numerical simulations and experimental results. In: Proceedings of the 33rd international Symposium on ballistics; 2023. Bruges.
[18] Held M. Shaped charge optimisation against ERA targets. Propellants, Explos Pyrotech 2005;30(3):216—23.
[19] Mickovic D, Jaramaz S, Elek P, Miloradovic N, Jaramaz D. A model for explosive reactive armor interaction with shaped charge jet. Propellants, Explos Pyrotech 2016;41(1):53—61.
[20] Bergstra J, Bardenet R, Bengio Y, Kegl B. Algorithms for hyperparameter optimization. In: Proceedings of advances in neural information processing systems (NeurIPS); 2011.
[21] Cully A, Clune J, Tarapore D, Mouret J-B. Robots that can adapt like animals. Nature 2015;521(7553):503—7.
[22] Barnett M, Senadeera M, Fabijanic D, Shamlaye KF, Joseph J, Kada SR, Rana S, Gupta S, Venkatesh S. A scrap-tolerant alloying concept based on high entropy alloys. Acta Mater 2020;200:735—44.
[23] Ryan S, Berk J, Rana S, McDonald B, Venkatesh S. A Bayesian optimisation methodology for the inverse derivation of viscoplasticity model constants in high strain-rate simulations. Defence Technology 2022;18(9):1563—77.
[24] Ryan S, Berk J, Weiss A, Yatom N, Vardy M, Rana S, Greenhill S, Venkatesh S. Adaptive experimental optimisation for sample-efficient armour design. In: Proceedings of the 32nd international Symposium on ballistics; 2022. Reno.
[25] Collé A, Limido J, Vila J-P. An accurate multi-regime SPH scheme for barotropic flows. J Comput Phys 2019;388:561—600.
[26] Moldtmann P. Application of Bayesian optimization for improving the design parameters of explosive reactive armor fit for a fixed number of different threats. Masters thesis. Helmut Schmidt University; 2023.
[27] Westerling L. Interaction of cylindrical penetrators with ceramic and electromagnetic armor. PhD thesis. Uppsala University; 2013.
[28] Skoglund P. Constitutive modelling and mechanical properties of a tungsten heavy metal alloy. Swedish Defence Research Agency (FOI), Tumba, FOI-R-0723—SE 2002.
[29] Collé A, Limido J, Unfer T, Vila J-P. Innovative meshless approach for shaped charges applications. J Phys Conf 2022;2154:012002.
[30] Borvik T, Hopperstad OS, Berstad T, Langseth M. A computational model of viscoplasticity and ductile damage for impact and penetration. Eur J Mech Solid 2001;20:685—712.
[31] Nilsson M. "Constitutive model for armox 500T and armox 600T at low and medium strain rate", Swedish defence research agency (FOI), tumba, FOI-R—1068—SE. 2003.
[32] Nguyen LH, Lässig TR, Ryan S, Riedel W, Mouritz AP, Orifici AC. A methodology for hydrocode analysis of ultra-high molecular weight polyethylene composite under ballistic impact. Composites Part A 2016;84:224—35.
[33] Juanicotena A. Theoretical and experimental study of high strain, high strain rate materials viscoplastic behaviour: application to Mars 190 steel and tantalum. PhD thesis Institut Supérieur de Génie Mécanique et Productique 1998.
[34] Helte A, Lundgren J, Örnhed H, Norrefeldt M. Evaluation of performance of m/46. FOI-R–2051–SE, Swedish Defence Research Agency; 2006 (in Swedish).
[35] Elbeih A, Hussein AK, Elshenawy T, Zeman S, Hammad SM, Baraka A, Elsayed MA, Gobara M, Tantawy H. Enhancing the explosive characteristics of

a Semtex explosive by involving admixtures of BCHMX and HMX. Defence Technology 2020;16:487–92.
[36] Rousseau DM. Schema, promise and mutuality: the building blocks of the psychological contract. J Occup Organ Psychol 2001;74(4):511–41.
[37] Mockus J. Application of Bayesian approach to numerical methods of global and stochastic optimization. J Global Optim 1994;4:347–65.
[38] Snoek J, Larochelle H, Adams RP. Practical Bayesian optimization of machine learning algorithms. 2012. *arXiv*:1206.2944.
[39] Liu H, Ong Y-S, Cai J. A survey of adaptive sampling for global metamodeling in support of simulation-based complex engineering design. Struct Multidiscip Optim 2018;57:393–416.
[40] Greenhill S, Rana S, Gupta S, Vellanki P, Venkatesh S. Bayesian optimization for adaptive experimental design: a review. IEEE Access 2020;8:13937–48.
[41] Li C, Rana S, Gupta S, Nguyen V, Venkatesh S, Sutti A, Rubin D, Slezak T, Height M, Mohammed M, Gibson I. Accelerating experimental design by incorporating experimenter hunches. In: Proceedings of the international conference on data mining (ICDM); 2018.
[42] Shilton A, Gupta S, Rana S, Venkatesh S. Regret bounds for transfer learning in Bayesian optimisation. Proceedings of Machine Learning Research (PMLR) 2017;54:307–15.
[43] Gupta S, Shilton A, Kumar Av A, Ryan S, Abdolshah M, Le H, Rana S, Berk J, Rashid M, Venkatesh S. BO-Muse: a human expert and AI teaming framework for accelerated experimental design. 2023. p. 01684v2. *arXiv*:2303.
[44] Daw ND, O'Doherty J, Dayan P, Seymour B, Dolan RJ. Cortical substrates for exploratory decisions in humans. Nature 2006;441(7095):876–9.
[45] Dane E. Reconsidering the trade-off between expertise and flexibility: a cognitive entrenchment perspective. Acad Manag Rev 2010;35(4):579–603.