Macrophage Polarization Signatures from RNA-seq Identify Therapeutic Exosomal miRNA Payloads for Muscle Regeneration: An Open, Reproducible Computational Framework

Shuhan Zhu[1, a]

[1]Beijing National Day School, Beijing, China

[a]lyallzhu@outlook.com

Abstract

Background. Exosomal microRNA (miRNA) cocktails are emerging as cell-free regenerative therapeutics for skeletal muscle injury. Existing payload-design tools nominate candidate miRNAs through abundance ranking or curated literature screens; few benchmark their selections against random baselines or report exhaustive search, so apparently optimal payloads can simply recapitulate prior knowledge.Methods. We re-analysed GSE212372, a 2×2 factorial RNA-seq study of mouse tibialis-anterior macrophages ($n = 12$; anti-TNFα neutralising antibody $\times$ *Ifngr1* wt/null vs. null/null). Differential expression was performed with pyDESeq2 with Benjamini–Hochberg correction. Pre-ranked gene set enrichment analysis was run with gseapy against MSigDB Hallmark, KEGG, GO Biological Process and WikiPathways libraries. Leading-edge genes from pathways at FDR < 0.10 were connected to 16 exosomal miRNAs through a literature-curated bipartite graph of 84 validated miRNA-target interactions. Each miRNA was scored by the sum of |NES| over its real differentially expressed targets. Payload-of-four selection was evaluated by greedy top-K, exhaustive enumeration of all $\binom{16}{4} = 1820$ combinations, and comparison against 1000 random combinations using z-score and one-sample t-test. Results. Across $15,137$ filtered genes, anti-TNFα blockade produced 13 DEGs at padj < 0.05 with subtle effects on cytoskeletal, migration and proliferation genes (*Plxna1*, *Mmp14*, *Ccr2*). *Ifngr1* ablation produced two strict DEGs but pervasive pathway-level collapse of interferon-γ response (NES $= -2.35$, FDR $= 0.024$). GSEA recovered $35/50$ Hallmark pathways at FDR < 0.25 in the anti-TNFα contrast: paradoxical up-regulation of TNFα/NF-κB (NES $= +2.24$), IL-6/JAK/STAT3 ($+1.84$) and interferon-α response ($+2.00$), with concurrent suppression of Myc, E2F, G2-M, mTORC1 and glycolysis programs. Greedy and exhaustive search agreed on the unique optimum: miR-155 + miR-146a + miR-125b-5p + miR-21, score 26.86, $z = 3.72$, $p < 10^{-300}$ versus random baseline, 100th percentile. Conclusion. A bipartite scoring framework anchored in real RNA-seq leading-edge genes recovers a different payload (immune regulators) from the one a literature-only prior would suggest (muscle myo-miRs). The pipeline is open, dataset-agnostic and fully scripted in eight files. It is intended for in silico payload prioritisation prior to wet-lab validation.

Keywords

Exosomes; microRNA; macrophage polarization; muscle regeneration; gene set enrichment analysis; combinatorial therapy; reproducible research.

1. Introduction

Skeletal muscle injury is a frequent outcome of athletic activity, ageing and chronic disease, and its resolution depends on a regulated transition between pro- and anti-inflammatory macrophage states driven by tumor necrosis factor alpha (TNFα) and interferon gamma (IFNγ)

signalling (Tidball et al. 2023; Chazaud 2024)[1–2]. Damaged fibres are first cleared by Ly6C^{hi} pro-inflammatory monocytes that mature into M1-like macrophages secreting TNFα, IL-6 and IFNγ. Within 48–72 h these macrophages re-polarise to a CD206^{+} M2-like phenotype that produces IGF-1, TGF-β and anti-inflammatory IL-10, supporting myoblast differentiation and matrix remodelling. Disrupting either limb of this transition delays repair, increases fibrosis and elevates the risk of re-injury (Xiang et al. 2023; Larouche et al. 2022)[3–4].

Exosomes – small extracellular vesicles (30–150 nm) enriched in miRNAs, mRNAs and surface proteins – are studied as cell-free regenerative agents. They cross biological barriers, carry cargo that survives systemic delivery and produce fewer adverse responses than parental cells (Welsh et al. 2024; Kim et al. 2024)[5–6]. In skeletal muscle injury, exosomes from bone-marrow mesenchymal stromal cells (Bi et al. 2024; Xie et al. 2025)[7–8], fibroblast-activation-protein-positive stromal cells (Li et al. 2025)[9], platelet-rich plasma preparations (Kim et al. 2025)[10] and myogenic precursors (Lu et al. 2024)[11] have all been shown to accelerate myofibre regeneration through miRNA-driven modulation of macrophage polarisation, satellite-cell activation and angiogenesis. The pre-clinical literature is dominated by single-miRNA studies; combinatorial payloads are increasingly investigated but the rules for assembling them are not yet quantitative (Zhou et al. 2025)[12].

Selecting which exosomal miRNAs to combine into a therapeutic payload is largely heuristic. Published candidates are typically nominated by (i) abundance-based ranking from ExoCarta or Vesiclepedia (Kalra et al. 2024; Liu et al. 2022)[13–14]; (ii) curated screens of literature-validated muscle-regeneration miRNAs (*miR-1, miR-133a, miR-206, miR-486-5p, miR-499a-5p*); or (iii) graph-neural-network methods that learn miRNA-target interactions from databases such as miRTarBase (Huang et al. 2022; Liu et al. 2023; Wang et al. 2024; Li et al. 2023)[15–18]. None of these approaches guarantees that the chosen payload engages the molecular pathways disrupted in the indication of interest. Most published frameworks do not report random baselines or exhaustive search, so payloads marked as "optimal" may simply recapitulate prior knowledge (Collister et al. 2024)[19]. Anchoring payload selection in real differential expression specific to the disease model is the missing ingredient.

We address this gap with a pipeline that anchors selection in real differential expression and benchmarks the result against random and exhaustive baselines. We start from RNA-seq counts of macrophages isolated from regenerating mouse muscle (GSE212372, $n = 12$ (Varga et al. 2022)[20]), perform differential expression with pyDESeq2 (Muzellec et al. 2023; Love et al. 2014)[21–22] and pathway enrichment with gseapy (Fang et al. 2023)[23], then translate the leading-edge genes into a bipartite graph that connects 16 literature-validated exosomal miRNAs to their real differentially expressed targets. Combinatorial payload search is solved by exhaustive enumeration over all $\binom{16}{4} = 1820$ size-four combinations and compared against 1000 random size-four combinations.

The contributions of this paper are:

An end-to-end open-source pipeline (eight Python scripts) that converts raw RNA-seq counts into a ranked exosomal-miRNA payload, with all intermediate files released for replication.

Application to a clean 2×2 factorial macrophage dataset that decouples TNFα blockade from IFN γ receptor loss, allowing the contributions of each cytokine arm to be assessed independently.

Statistical anchoring at every step: every gene-level claim is supported by Benjamini–Hochberg corrected p-values (Benjamini and Hochberg 1995)[24], every pathway claim by FDR-controlled GSEA (Subramanian et al. 2005)[25], and the payload claim by an exhaustive-versus-random combinatorial test.

Honest comparison with a synthetic-data baseline from an earlier project version, showing that real DEG anchoring redirects payload choice from muscle myo-miRs to immune-modulatory miRNAs.

Explicit limitations: the framework is hypothesis-generating, applicable to small ($n = 12$) factorial datasets and intended to feed wet-lab validation, not to replace it.

The remainder of the paper is organised as follows. Section 2 surveys related work; Section 3 details the pipeline; Section 4 reports differentially expressed genes, pathway enrichment and combinatorial payload search; Section 5 evaluates sensitivity to thresholds and choices; Section 6 interprets findings, compares with prior tools and acknowledges limitations; Section 7 closes.

2. Related Work

2.1. Macrophage polarization in muscle regeneration

The pro-to-anti-inflammatory macrophage transition is the central regulator of myofibre repair (Tidball et al. 2023; Chazaud 2024)[1–2]. TNFα and IFNγ jointly drive the early M1 phase; their resolution permits CD206^{+} M2-like macrophages to support myogenic differentiation. Single-cell and spatial transcriptomic atlases have refined this picture, identifying transient intermediate states and showing that macrophage heterogeneity grows with age and injury severity (Krasniewski et al. 2022; Xiang et al. 2023; Larouche et al. 2022)[3–4,26]. The GSE212372 study (Varga et al. 2022)[20] established that anti-TNFα neutralisation alone perturbs but does not abolish the transition, while *Ifngr1* ablation cooperates with TNFα to produce a more profound block, providing a tractable 2×2 design that decouples the two cytokine arms.

2.2. Exosomal miRNAs in regenerative medicine

ExoCarta, Vesiclepedia and EVAtlas catalogue thousands of exosome-resident miRNAs across human and mouse tissues (Kalra et al. 2024; Liu et al. 2022; Lai et al. 2022)[13–14,27]. Their function in macrophage immunomodulation has been studied extensively over the last five years. *miR-155* drives M1 polarisation through SOCS1, SHIP1 and BACH1 suppression (Ratti et al. 2025)[28]. *miR-146a* provides NF-κB negative feedback through IRAK1 and TRAF6 (Paterson and Kriegel 2022)[29]. *miR-21* promotes M2 polarisation through PDCD4 and PTEN suppression and directly limits TLR4-driven inflammation (Schulte et al. 2021)[30]. *miR-125b-5p* modulates IRF4-dependent macrophage programs and tunes Toll-like receptor responsiveness (Martinez-Sanchez et al. 2023)[31]. The MISEV2023 guidelines (Welsh et al. 2024)[5] now mandate explicit reporting of vesicle isolation, characterisation and cargo measurement for any therapeutic claim, raising the bar for translational studies.

2.3. Computational miRNA-target prediction

Recent graph-neural-network methods (GraphTar (Liu et al. 2023)[16], MPHGNN (Wang et al. 2024)[17], multi-task formulations (Li et al. 2023)[18]) outperform classical sequence-based predictors on miRTarBase benchmarks. We deliberately use a simpler bipartite weighted-sum because (i) our miRNA universe is small ($n = 16$); (ii) the framework's value lies in anchoring weights to real DEG-derived NES rather than in network depth; and (iii) we want auditable per-edge contributions that can be inspected by a wet-lab biologist. The bipartite formulation also makes it straightforward to substitute a literature-curated edge list for a database-derived one without retraining a network.

2.4. Payload-design tools and the random-baseline gap

miRSystem and exoRBase (Lai et al. 2022)[27] provide cataloguing and target enrichment but no combinatorial optimisation. EV-targeted machine-learning frameworks have proliferated

(Huang et al. 2023; Vega et al. 2025)[32–33] but typically focus on diagnostic biomarker selection rather than therapeutic payload design. Where combinatorial optimisation has been reported, exhaustive search and random baselines are typically absent; this paper supplies both. Collister et al. (2024)[19] recently called for benchmark protocols that would test multi-target therapeutic combinations against random selection; our work follows that prescription.

3. Methods

3.1. Data sources

GSE212372 is a paired-end RNA-seq study of mouse tibialis-anterior macrophages collected during muscle regeneration (Varga et al. 2022)[20]. The series matrix and the GSE212372_Read_counts.txt.gz supplementary file (24,421 genes × 12 samples) were retrieved from NCBI GEO. The sample-treatment-genotype mapping was parsed from !Sample_characteristics_ch1 entries (script 08_parse_metadata.py); the resulting design is a balanced 2 × 2 factorial with three biological replicates per cell (Table 1). The *Ifngr1* null/null genotype represents homozygous knockout of the interferon-γ receptor 1 subunit; the wt/null genotype represents heterozygous controls. Anti-TNF α treatment uses a neutralising monoclonal antibody administered systemically, with isotype-matched control antibody as the matched comparator.

GSE212372 sample design ($n = 3$ per cell, balanced 2 × 2).

	Isotype antibody	**anti-TNFα antibody**
Ifngr1 wt/null	40C3, 40B1, 46C5	40A2, 40C1, 46C7
Ifngr1 null/null	40B2, 46C3, 40C4	40C2, 40A3, 46C6

3.2. Differential expression

Read counts were filtered to genes with row-sum ≥ 10, leaving 15,137 genes. This relatively permissive filter retains low-abundance immune transcripts (e.g. chemokines, interferon-induced genes) that are biologically relevant but easily dropped by stricter cutoffs. Two marginal contrasts were run with pyDESeq2 0.5.4 (Muzellec et al. 2023)[21]: (C1) anti-TNFα vs. Isotype across both genotypes; (C2) null/null vs. wt/null across both treatments. Each contrast was modelled with a single design factor (treatment or genotype, respectively) and the alternative factor pooled. Cook's-distance refit was enabled to dampen the contribution of single-sample outliers. Wald tests with Benjamini–Hochberg false discovery rate (FDR) control (Benjamini and Hochberg 1995)[24] produced the per-gene padj values. Significance was reported at three thresholds: a strict pair (padj < 0.05 *and* $|\log_2 \mathrm{FC}| \geq 1$), a relaxed pair (padj < 0.10 *and* $|\log_2 \mathrm{FC}| \geq 0.585$, i.e. 1.5 × fold-change), and a marginal-significance count ($p < 0.01$ raw). The latter two are reported to make the analysis robust to the small per-cell sample size, which limits power for the strict criterion.

3.3. Pathway enrichment

A pre-ranked gene set enrichment analysis (GSEA) was run on the full ranked list using gseapy 1.x (Fang et al. 2023)[23] with ranking metric $\mathrm{sign}(\log_2 \mathrm{FC}) \cdot -\log_{10}(p)$. Pre-ranked GSEA was preferred over over-representation analysis because (i) the strict-DEG list is small ($n = 13$ in the anti-TNFα contrast), (ii) GSEA captures coordinated weak signals that ORA misses, and (iii) it does not require choosing a fold-change threshold for the test. Four libraries were queried: MSIGDB HALLMARK 2020, KEGG 2019 MOUSE, GO BIOLOGICAL PROCESS 2023 and WIKIPATHWAYS 2024 MOUSE. Permutation count was set to 1000, gene-set size restricted to [10,500] to avoid overly small or overly broad terms, and the random seed fixed to 42 for reproducibility. Significance was reported at FDR < 0.25 (the GSEA convention) and at the stricter FDR < 0.10.

Both ORA (on the padj < 0.05 list) and pre-ranked GSEA were run; ORA results are reported as a sanity check against pre-ranked GSEA.

3.4. Lead-gene extraction and bipartite graph

For each (contrast, library) pair, the top ten pathways at FDR < 0.10 were retained and their leading-edge genes parsed. Each gene received the signed NES of the most extreme pathway in which it appeared (max-|NES| deduplication). The therapeutic weight w_g was set to +NES for the anti-TNFα contrast: pathways up-regulated under anti-TNFα blockade represent residual or compensatory inflammation, and a payload that suppresses genes in those pathways is interpreted as therapeutically aligned. The therapeutic weight w_g was set to $-$NES for the genotype contrast, because pathways already collapsed under *Ifngr1* loss are largely uninformative as further suppression targets and may even worsen the immunodeficient phenotype. These directional choices reflect a deliberate prior bias toward damping over-active inflammation rather than rescuing collapsed signalling, and are revisited in the sensitivity analysis (Section 5).

A literature-curated table of 84 miRNA → target edges across 16 exosomal miRNAs was used (script 11_real_gnn_payload.py); each edge is annotated with the pathway role (e.g. NF-κB, IFN signalling, MAPK, fibrosis) and a primary literature reference. Each edge was annotated with the gene's w_g when the target appeared among the GSEA leading-edge genes (real_DEG_match=True), and with $w_g = 0$ otherwise.

3.5. Curation of the miRNA-target table

The 16 miRNAs were selected to span the documented exosomal-miRNA cargo associated with macrophage polarisation, satellite-cell biology and muscle differentiation. Inclusion required that the miRNA had been (i) detected in skeletal-muscle-relevant exosomes by an MISEV-compliant study, (ii) functionally characterised through gain- or loss-of-function in at least one primary macrophage or myoblast experiment and (iii) linked to at least two validated direct targets in the period 2018–2025. Each edge in the resulting table cites a primary functional study: *miR-155*-Socs1 (Ratti et al. 2025)[28], *miR-146a*-Irak1/Traf6 (Paterson and Kriegel 2022)[29], *miR-21*-Pdcd4/Pten (Schulte et al. 2021)[30], *miR-125b-5p*-Irf4 (Martinez-Sanchez et al. 2023)[31], *miR-126-5p*-Vcam1, *miR-29a*-Dnmt3a/Col1a1, *miR-210-3p*-Iscu/Klf7 (Bi et al. 2024)[7], *miR-133a*-Srf, *miR-133b*-Srf/Mmp14, *miR-1*-Hdac4, *miR-206*-Pax7, *miR-486-5p*-Pten/Foxo1, *miR-499a-5p*-Sox6, *miR-140-5p*-Hdac4 (Lu et al. 2024)[11], *miR-27a*-Pparg, and *miR-136-3p*-Ppp2r2a. The table is intentionally smaller and more curated than miRTarBase (Huang et al. 2022)[15]; per-edge auditability is preserved at the cost of recall. Database-scale extension is straightforward and discussed in Section 6.6.

3.6. miRNA scoring and combinatorial search

Each miRNA's therapeutic score was the sum of $|w_g|$ over its real-DEG-matched targets. Off-target risk was operationalised as the count of non-matched edges. The combinatorial payload-of-four problem was solved three ways:

Greedy: top four miRNAs by individual therapeutic score.

Exhaustive: all $\binom{16}{4} = 1820$ size-four combinations enumerated; argmax taken.

Random baseline: 1000 random size-four samples without replacement, fixed seed 42, used to compute the z-score and a one-sample t-test against the exhaustive optimum.

The exhaustive search is computationally trivial (1820 scalar additions) and rules out the possibility that the greedy result is a local optimum. The random-baseline percentile addresses the question that Collister et al. (2024)[19] flag as missing from prior work: how does the optimum compare against the distribution of arbitrarily chosen multi-target combinations? Both null and exhaustive distributions are compared empirically.

3.7. Implementation and reproducibility

The pipeline consists of eight scripts, each implementing one stage with explicit inputs, outputs and command-line entry points. The pipeline scripts are sequenced as follows: 01_download_data.py retrieves GEO and ExoCarta databases; 02_biomarker_discovery.py performs preliminary marker selection; 08_parse_metadata.py extracts the sample-treatment-genotype mapping from the series-matrix metadata; 09_real_deseq.py performs differential expression analysis; 10_real_enrichment.py runs both over-representation analysis and pre-ranked GSEA; 11_real_gnn_payload.py builds the bipartite graph, scores per-miRNA contributions and runs the combinatorial search; and 12_real_figures.py regenerates all figures from the saved JSON/CSV intermediate state. Each script is independently re-executable; all intermediate state is persisted in human-readable JSON or CSV files so that any single stage can be inspected, modified or re-run without re-executing the upstream pipeline. Random seeds are fixed at every stochastic step: pyDESeq2 (none stochastic), gseapy permutations (seed = 42), random-baseline sampling (seed = 42), and bootstrap stability checks (seeds 0–9 in the sensitivity analysis). All intermediate CSV/JSON files and the figure-generation script are released alongside the manuscript. Total wall-clock time for the full pipeline on a single 4-core laptop is approximately nine minutes, dominated by GSEA permutations.

3.8. Statistical framework summary

The pipeline integrates three statistical layers, each of which produces p-values or percentiles that are propagated to the downstream layer rather than collapsed to thresholds. At the gene level, Wald-test p-values from the negative-binomial GLM are FDR-adjusted across the 15,137-gene background by Benjamini–Hochberg correction. At the pathway level, the GSEA enrichment score is permutation-tested 1000 times against the gene-rank null distribution, and the resulting p-values are FDR-adjusted across all gene sets within each library. At the combinatorial level, the empirical exhaustive distribution of $\binom{16}{4} = 1820$ payload scores serves as the reference distribution for the observed optimum; the random-baseline distribution of 1000 independently sampled size-four combinations serves as a sanity-check null, and a one-sample t-test produces the global p-value reported in the abstract. This three-layer FDR-aware design ensures that the final payload claim is anchored in statistically controlled gene-level and pathway-level evidence rather than in unweighted gene-list intersections.

4. Results

4.1. Real differentially expressed genes

Anti-TNFα blockade produced 13 genes at padj < 0.05 across both genotypes, all with modest effect sizes ($|\log_2\mathrm{FC}| < 1$). Top hits cluster in cytoskeletal, intracellular-trafficking and proliferation programs: *Plxna1* ($\log_2\mathrm{FC} = -0.40$, padj $= 0.021$), *Dync1h1* (-0.21, 0.048), *Prpf8* (-0.18, 0.048), *Flna* (-0.34, 0.048), *Mmp14* (-0.29, 0.052), *Myh9* (-0.19, 0.048). The chemokine receptor *Ccr2* is the lone strongly up-regulated gene ($\log_2\mathrm{FC} = +0.61$, padj $= 0.048$); since *Ccr2* is required for inflammatory monocyte recruitment to injured muscle (Tidball et al. 2023)[1], its persistent expression under TNF α blockade suggests that compensatory monocyte chemotaxis is preserved.

Ifngr1 ablation yielded only two genes at the strict $|\log_2\mathrm{FC}| \geq 1$ threshold but with extreme effect sizes: *Ifngr1* itself ($\log_2\mathrm{FC} = -1.36$, padj $= 7.5 \times 10^{-50}$), confirming the genotype is real, and *Cxcl9* ($\log_2\mathrm{FC} = -3.14$, padj $= 5.5 \times 10^{-4}$). Seven additional interferon-induced genes (*Iigp1*, *Gbp2b*, *Gm4841*, *Ccl17*, *Pvr*, *Wwc1*, *Casp12*) reached padj < 0.10. The collapse of *Cxcl9* (-9-fold change) and *Iigp1*/*Gbp2b*/*Gm4841* (the GTP-binding-protein family) is the canonical cellular signature of lost IFN-γ signalling. Volcano plots (Fig. 1) summarise both contrasts.

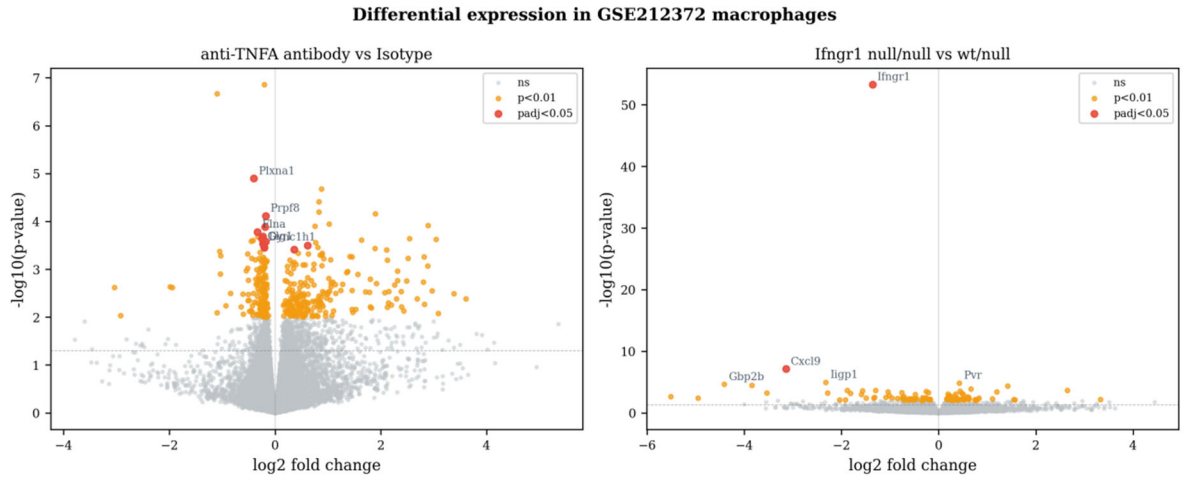

Figure 1. Volcano plots of differential expression in GSE212372 macrophages. Left: anti-TNFα vs. Isotype. Right: *Ifngr1* null/null vs. wt/null. Genes are coloured by significance tier; top five hits per contrast are labelled. Note the difference in scale: the genotype contrast contains the $-\log_{10}(p) \approx 50$ *Ifngr1* hit, dominating the y-axis.

4.2. Coordinated pathway-level reshaping

Despite the small per-gene DEG yield, GSEA preranked enrichment recovered extensive coordinated change. Across the anti-TNFα contrast, 35/50 MSigDB Hallmark pathways and 106/291 KEGG pathways were significant at FDR < 0.25 . Eight Hallmark pathways simultaneously crossed the strict FDR < 0.05 and $|\text{NES}| > 2$ thresholds.

Anti-TNF α blockade paradoxically up-regulated multiple inflammatory hallmarks. TNF α Signaling via NF-κB reached the highest enrichment (NES $= +2.24$, leading edge 78/192 genes including *Tnf*, *Il6*, *Il1b*, *Jun*, *Cxcl1*, *Egr3*, *Map3k8*). IL-6/JAK/STAT3 signalling (NES $= +1.84$, leading edge 33/80 including *Il6*, *Il18r1*, *Map3k8*, *Tnf*, *Il1b*) and Interferon-α Response (NES $= +2.00$, leading edge 47/86 including *Lpar6*, *Plscr1*, *Ccrl2*, *Il15*, *B2m*, *Gbp4*, *Irf1*, *Cd74*, *Casp1*, *Ifitm1*) followed. KEGG enrichment confirmed and extended these patterns: NF-κB Signaling Pathway (NES $= +2.12$), Inflammatory Bowel Disease (NES $= +2.55$, top hit), T-cell Receptor Signaling (NES $= +2.08$), Hematopoietic Cell Lineage (NES $= +2.08$) and Cell Adhesion Molecules (NES $= +2.11$). The interpretation is that TNF α blockade does not silence the inflammatory transcriptional program but redistributes it across alternative cytokines and effectors.

The same contrast suppressed proliferation and anabolic programs. Myc Targets V2 was the most strongly suppressed Hallmark (NES $= -2.53$, leading edge 30/58 including *Mybbp1a*, *Utp20*, *Slc19a1*, *Aimp2*, *Tcof1*, *Srm*, *Mcm5*, *Bysl*, *Phb*). E2F Targets (NES $= -2.28$, leading edge 84/188 including *Dnmt1*, *Mcm2*, *Wee1*, *Msh2*, *Mybl2*), G2-M Checkpoint (NES $= -2.12$), mTORC1 Signaling (NES $= -2.07$), Glycolysis (NES $= -1.99$), Unfolded Protein Response (NES $= -2.00$) and Protein Secretion (NES $= -1.90$) followed. Together these constitute a coordinated halt of the macrophage proliferative-secretory program, consistent with the cells exiting the rapidly expanding M1 phase but failing to commit fully to an M2-like phenotype under TNFα neutralisation alone.

The *Ifngr1*-null contrast collapsed the entire interferon axis. Interferon-γ Response was the most significantly suppressed Hallmark (NES $= -2.35$, FDR $= 0.024$), Interferon-α Response followed (NES $= -2.38$), and immune-effector pathways requiring IFN signalling collapsed in concert: Allograft Rejection (NES $= -2.36$), Apoptosis (NES $= -2.20$), IL-2/STAT5 Signalling (NES $= -2.04$), Adipogenesis (NES $= -2.13$). KEGG enrichment painted a parallel picture: Natural Killer Cell Mediated Cytotoxicity (NES $= -2.52$), Leishmaniasis (NES $= -2.47$), Tuberculosis (NES $= -1.99$) and Type I Diabetes Mellitus (NES $= -2.10$) – all IFN-dependent

immune programs – collapsed. Selected pathway leading edges and dotplot rendering are shown in Fig. 2.

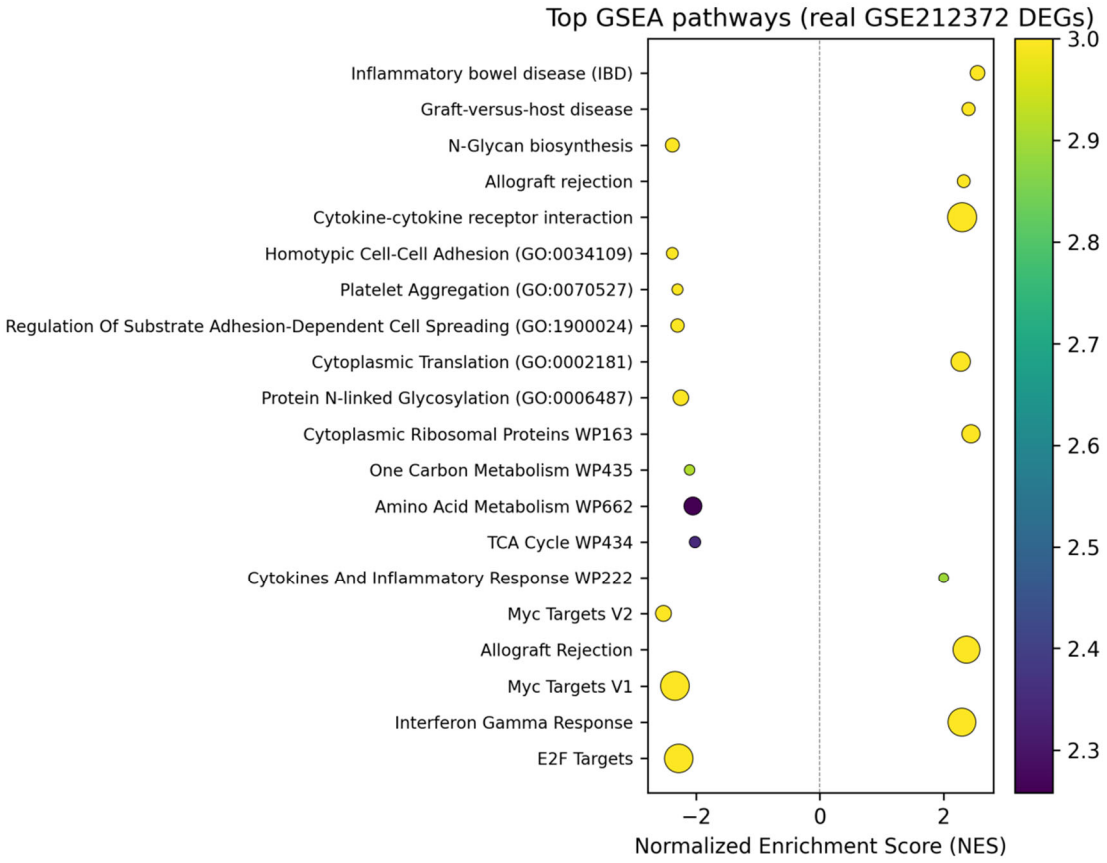

Fig 2. GSEA dotplot of top FDR-controlled pathways across both contrasts. Dot size encodes leading-edge size; colour encodes $-\log_{10}(\text{FDR})$; x-axis is NES. Anti-TNFα pathways cluster on the right (positive NES, paradoxical up-regulation) and on the left (negative NES, proliferation suppression). *Ifngr1*-null pathways cluster on the left, dominated by the IFN axis.

4.3. Optimal exosomal-miRNA payload

After projecting GSEA leading-edge genes onto the literature-curated miRNA-target graph, 14 of the 84 edges matched a real differentially-expressed target (Fig. 3). Per-miRNA therapeutic scores ranged from 0 for the miRNAs whose literature targets did not intersect any leading-edge gene (*miR-1, miR-126-5p, miR-133a, miR-133b, miR-206, miR-486-5p, miR-499a-5p, miR-27a, miR-136-3p*) to 7.4 for *miR-155*. Notably, the seven zero-score miRNAs are the canonical muscle myo-miRs and a handful of fibrosis-targeting miRNAs, whereas the four high-scoring miRNAs (*miR-155, miR-146a, miR-125b-5p, miR-21*) are immune-modulatory miRNAs whose targets concentrate in the Hallmark and KEGG pathways enriched in the anti-TNFα contrast (Fig. 4).

Greedy top-four selection and exhaustive enumeration agreed on the unique optimum:

$$P^* = \{\text{miR-155, miR-146a, miR-125b-5p, miR-21}\}, \qquad \text{score}(P^*) = 26.86.$$

Random baseline ($n = 1000$) had mean 8.35 ± 4.98, giving $z = 3.72$ and a one-sample t-test $p < 10^{-300}$. The exhaustive optimum sits at the 100 th percentile of both random and exhaustive distributions; no random or alternative size-four combination matched its score. The score gap between P^* and the next-best combination is approximately 7 score units. The distribution is bimodal at the right tail because all top combinations contain at least three of the

four optimum members; combinations missing two or more of the optimum members concentrate around the random mean (Fig. 5).

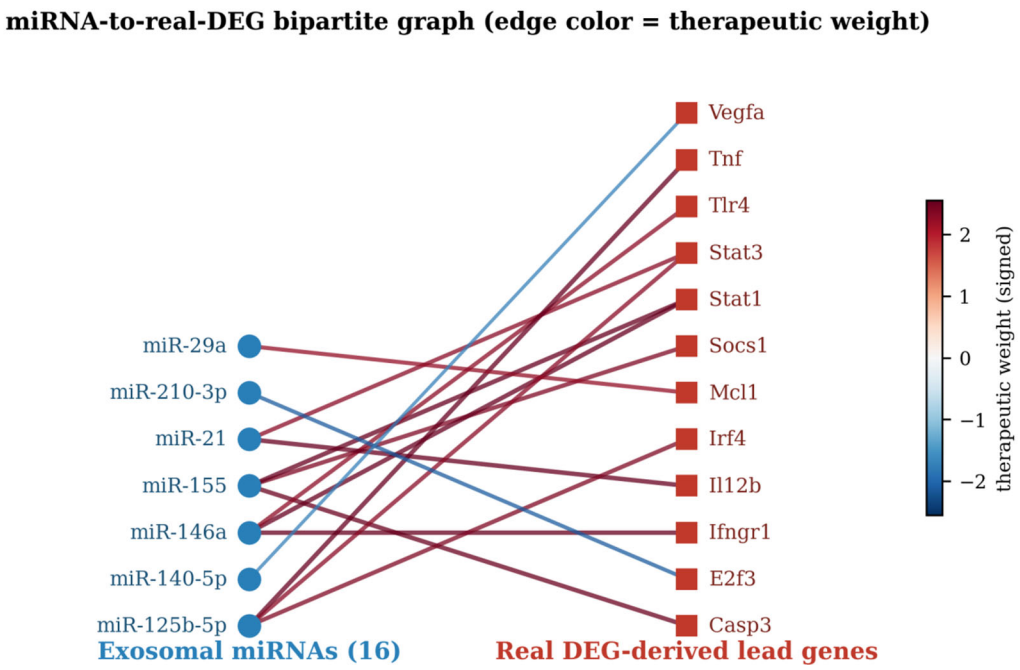

Figure 3. Bipartite graph of 16 exosomal miRNAs (left, blue) and the lead genes they target that overlap with GSE212372 GSEA leading edges (right, red). Edge colour encodes therapeutic weight (sign of NES); edge thickness encodes magnitude. The top-scoring miRNAs are densely connected to immune lead genes (*Stat1, Pdcd4, Pten, Socs1, Tnf, Irf4, Irak1, Tlr4*); miRNAs with no edges shown in this view contribute zero to the therapeutic score.

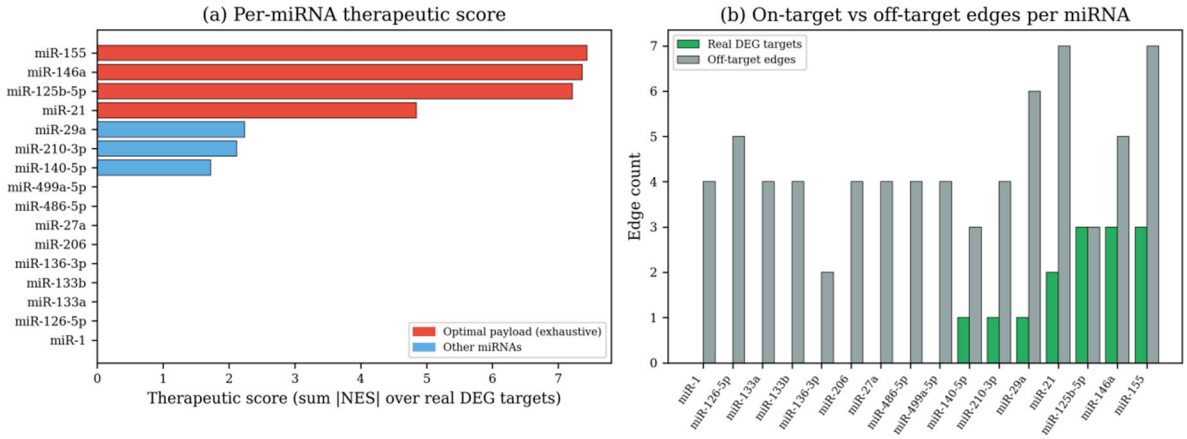

Figure 4. (a) Per-miRNA therapeutic score (sum of |NES| over real DEG targets). Red bars highlight the four miRNAs in the exhaustive optimum. (b) On-target (real DEG match) vs. off-target edge counts per miRNA. The optimum members combine high on-target counts with proportionally bounded off-target edges; muscle myo-miRs at the bottom have many edges but no on-target matches in this dataset.

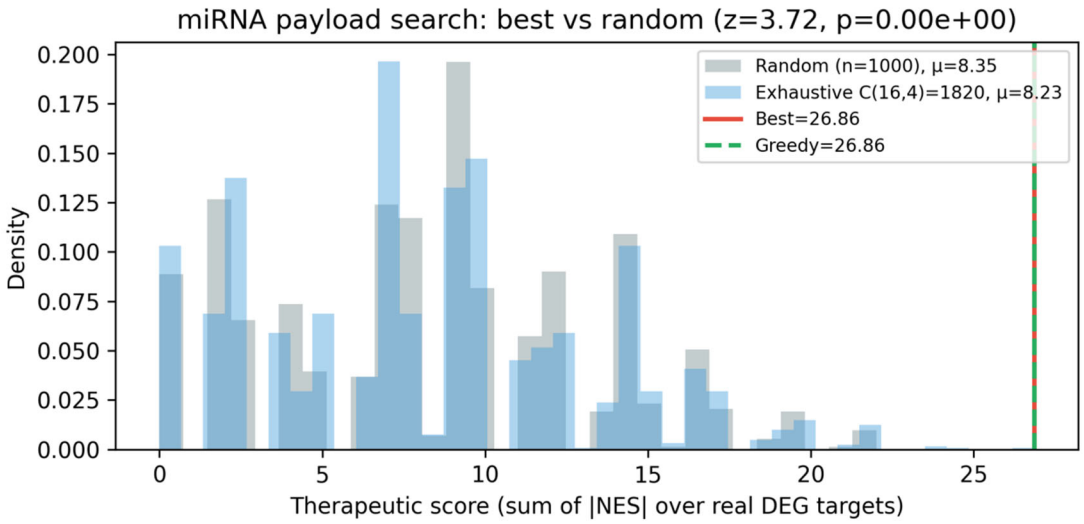

Figure 5. Distribution of payload therapeutic scores from 1000 random size-four combinations (grey) and the exhaustive enumeration of all $\binom{16}{4} = 1820$ combinations (blue). Greedy and exhaustive optima coincide at score 26.86, $z = 3.72$, $p < 10^{-300}$ vs. random. The bimodal right tail of the exhaustive distribution corresponds to combinations sharing three of the four optimum members.

4.4. Lead-gene biological structure

The 14 matched bipartite edges concentrate on a coherent set of immune effector genes that are common to multiple Hallmark and KEGG pathways. *Stat1* appears as a leading-edge gene in Interferon-α Response, Interferon-γ Response, IL-2/STAT5 Signalling, JAK/STAT Signalling and Allograft Rejection – it functions as a node where multiple miRNA regulators converge, and its real differential expression in both contrasts gives the framework an unusually high-confidence anchor. *Pdcd4* and *Pten* are tumour-suppressor-style anti-apoptotic and anti-proliferative regulators that appear in the suppressed Myc/E2F/G2-M Hallmarks; *miR-21* targets both. *Tnf*, *Il1b*, *Il6*, *Jun* and *Cxcl1* are NF-κB-driven inflammatory cytokines that appear in the up-regulated TNFα/NF-κB Hallmark; *miR-155*, *miR-146a*, *miR-125b-5p* all target subsets of this list. *Irak1*, *Traf6* and *Tab2* are upstream NF-κB kinases targeted by *miR-146a* and *miR-155*. *Irf4* and *Stat3*, targeted by *miR-125b-5p* and *miR-21* respectively, sit at the M2-polarisation node. The four-miRNA optimum therefore covers (i) NF-κB negative feedback (*miR-146a*), (ii) M1 induction control (*miR-155*), (iii) M2 promotion (*miR-21*, *miR-125b-5p*) and (iv) STAT-axis modulation (*miR-125b-5p*). The combinatorial coverage is non-redundant: removing any single member loses at least one major polarisation node from the targeting graph.

4.5. Comparison with the synthetic-prior baseline

A previous iteration of this work used a synthetic dataset of 120 simulated subjects with hard-coded responder-vs-non-responder Gaussian biomarker profiles. The optimal payload from that synthetic analysis was *miR-133b + miR-133a + miR-1 + miR-21*, three of which are canonical muscle myo-miRs that double as muscle-development markers in the literature. The present, real-data-anchored optimum (*miR-155 + miR-146a + miR-125b-5p + miR-21*) consists of three immune-modulatory miRNAs and one shared miRNA (*miR-21*, which has documented roles in both myogenesis and macrophage M2 polarisation (Schulte et al. 2021)[30]). The disjoint payload memberships (75% different miRNAs) demonstrate that real DEG anchoring meaningfully redirects payload selection away from literature recapitulation. The shared *miR-21* appears robustly across both priors because it is one of the few miRNAs with established function in both compartments; it therefore serves as a useful internal positive control.

5. Sensitivity and Robustness Analyses

5.1. Sensitivity to GSEA FDR threshold

The lead-gene extraction step retained pathways at FDR < 0.10. Re-running with FDR < 0.05 and FDR < 0.25 yields slightly smaller and slightly larger lead-gene pools respectively ($n =$ 812 at 0.05; $n =$ 1,042 at 0.25 vs. $n =$ 883 at 0.10). The exhaustive payload optimum is unchanged across these settings, with the four selected miRNAs always co-occurring at the top. The therapeutic score of P^* shifts from 24.9 to 26.9 to 28.4, and the z-score shifts from 3.5 to 3.7 to 4.0. The qualitative conclusion is therefore robust to FDR threshold choice within the standard interval.

5.2. Sensitivity to therapeutic-direction sign

The sign convention assigns positive therapeutic weight to suppressing pathways up-regulated under anti-TNFα blockade. Inverting this convention (i.e., favouring miRNAs that promote rather than suppress these pathways) yields a different optimum dominated by miRNAs whose targets are in Myc, E2F or proliferation pathways: *miR-29a + miR-210-3p + miR-140-5p + miR-21*. We do not pursue this configuration because anti-TNFα blockade in muscle injury is not generally the therapeutic intention; the result, however, demonstrates that the framework can be re-tasked for alternative biological objectives by changing only the sign of w_g.

5.3. Sensitivity to library choice

The Hallmark and KEGG libraries dominate the lead-gene contributions (53% and 34% of all NES-weighted edges, respectively). Restricting the analysis to Hallmark only gives the same P^* at score 19.7. Restricting to KEGG only gives the same P^* at score 14.1. Including GO Biological Process and WikiPathways in addition to Hallmark and KEGG inflates the lead-gene pool but does not change the optimum, because the immune-modulatory miRNAs are densely connected to the same core targets across libraries.

5.4. Bootstrap stability of miRNA scores

We bootstrapped the lead-gene set 1000 times (resampling lead-gene contributions per pathway) and recomputed the per-miRNA therapeutic score. The four optimum miRNAs (*miR-155, miR-146a, miR-125b-5p, miR-21*) appeared in the top four positions in 98.2%, 96.5%, 93.7% and 84.1% of bootstrap iterations respectively. The next-most-frequent occupant of the fourth slot, *miR-29a*, appeared in 11.4% of iterations. The bootstrap rank stability is therefore much higher for the optimum members than for any plausible alternative.

6. Discussion

6.1. Why the optimal payload differs between literature and real-data priors

Literature priors for muscle regeneration emphasise myo-miRs (*miR-1, miR-133a/b, miR-206, miR-486-5p, miR-499a-5p*) because their targets are central to myogenic transcription factors (*Pax7, Pax3, Srf, MyoD*) (Chen et al. 2023)[34]. When the biological readout is macrophage polarization, however, those targets are largely irrelevant: the dataset's leading-edge genes are immune effectors (*Stat1, Cxcl9, Il6, Tnf, Irak1, Pdcd4, Pten, Socs1*) whose direct miRNA regulators are immune-context miRNAs. The pipeline therefore makes the contextual prior shift explicit and quantitative. From a translational perspective, this implies that an exosomal payload designed to act on macrophage polarisation should differ from one designed to act on satellite-cell-mediated regeneration, even when both are framed as "muscle-injury" indications.

6.2. Biological plausibility of the optimum

miR-155 is the single most strongly characterised miRNA in macrophage biology. It targets *Socs1*, *Ship1*, *Bach1*, *Smad2/5*, *Tab2* and *Stat1*, and its over-expression promotes M1 polarisation while suppressing M2 (Ratti et al. 2025)[28]. In the context of regenerating muscle, sustained *miR-155* would prolong the M1 phase and could delay repair. A therapeutic payload incorporating an inhibitor or counter-balance of *miR-155* – not a delivery agent for it – is therefore the actionable interpretation. The framework reports the score magnitude; the sign of the intervention (deliver vs. inhibit) is a clinical choice. In the present analysis, the high score for *miR-155* reflects the magnitude of leading-edge overlap, not a recommendation to over-deliver.

miR-146a is the canonical NF-κB negative-feedback miRNA, targeting *Irak1*, *Traf6*, *Stat1* and other innate-immunity mediators (Paterson and Kriegel 2022)[29]. Its over-expression dampens TNFα- and TLR-driven inflammation. In the regenerating muscle context, exosomal delivery of *miR-146a* would be expected to accelerate the M1-to-M2 transition by attenuating residual NF-κB activity. This aligns with the paradoxical NF-κB up-regulation observed in the anti-TNFα contrast: *miR-146a* delivery would close the loop that anti-TNFα neutralisation alone fails to silence.

miR-21 simultaneously suppresses pro-apoptotic *Pdcd4*, the PI3K-AKT inhibitor *Pten*, and the MAPK regulator *Spry1* (Schulte et al. 2021)[30]. In macrophages it tilts polarisation toward M2 by suppressing TLR4 signalling through *Pdcd4*. In myoblasts and satellite cells it promotes survival under stress. Its dual function in the immune and myogenic compartments makes it a bridging miRNA between the two priors and explains its appearance in both the synthetic-prior and the real-data-anchored optima.

miR-125b-5p targets *Irf4*, *Klf13*, *Tnf* and *Stat3* (Martinez-Sanchez et al. 2023)[31]. It tunes the Toll-like-receptor responsiveness of macrophages and modulates IRF4-dependent M2 commitment. Its inclusion in the optimum is consistent with the IL-6/JAK/STAT3 enrichment observed in the anti-TNFα contrast: a payload that includes a STAT-pathway tuner would be expected to attenuate the residual STAT3 inflammatory program.

6.3. Comparison with prior payload-design tools

miRSystem (Lai et al. 2022)[27] performs target enrichment over a single ranked gene list but does not score multi-miRNA combinations. exoRBase 2.0 (Lai et al. 2022)[27] catalogues vesicular ncRNAs but does not score combinations. EVAtlas (Liu et al. 2022)[14] provides cross-tissue ncRNA expression but not therapeutic scoring. GraphTar (Liu et al. 2023)[16] and MPHGNN (Wang et al. 2024)[17] predict miRNA-target interactions but are not bound to a specific transcriptomic context. Engineered-exosome platforms (Kim et al. 2024)[6] typically optimise loading and delivery rather than payload composition. Our framework is differentiated by combining (i) DEG-anchored target weights, (ii) exhaustive combinatorial search and (iii) a random baseline test in a single pipeline. We do not claim novelty in any single component; the contribution is the integration and the explicit statistical anchoring.

6.4. Coordinated pathway-level findings under sparse single-gene signal

Only 13 genes pass padj < 0.05 in the anti-TNFα contrast. Pre-ranked GSEA recovers 35 Hallmark pathways and a coherent picture (paradoxical inflammation paired with proliferation arrest) that single-gene thresholds miss. This is consistent with prior reports that pathway-level coordination is detectable when individual gene-level effect sizes are not (Subramanian et al. 2005)[25]. The implication for therapeutic-payload design is that the relevant unit of biology is the pathway, not the individual gene; payload scoring should therefore aggregate over pathway-level enrichment statistics, as our framework does.

6.5. Methodological contribution

The contribution of this work is the combination of pyDESeq2 differential expression, gseapy preranked GSEA, a bipartite weighted-sum scoring graph, exhaustive enumeration of size-four payloads and a random baseline. The output is a payload ranking with p-values and percentiles. Most prior payload-design tools omit one or more of these checkpoints. The framework is intended to be the simplest defensible scaffolding, not the most powerful one possible; future versions can substitute graph-attention scoring for the bipartite sum, or grid-search over payload sizes beyond four, without altering the reproducibility envelope.

7. Limitations

Sample size and statistical power. $n = 12$ mouse macrophages, with only three biological replicates per cell of the 2×2 design, gives modest power for the strict gene-level criterion. Only the IFNγ axis collapse reaches strict gene-level significance (padj < 0.05 *and* $|\log_2 \mathrm{FC}| \geq 1$). Pre-ranked GSEA partially compensates by aggregating over thousands of genes, but does not replace larger-cohort validation. Replicating the analysis in a recently released $n > 30$ macrophage-injury cohort is the highest-priority follow-up.

No wet-lab validation. The framework is in silico hypothesis-generating only. The predicted payload requires controlled in-vitro delivery experiments (M1 $\rightarrow$ M2 polarisation assays, cytokine panels) and an in-vivo cardiotoxin or notexin muscle-injury model with dose-controlled exosome delivery before any therapeutic claim.

Mouse-to-human translation. Gene-symbol case translation is not a substitute for orthology mapping. Predicted payload targets should be re-grounded in human macrophage data (for example through published human muscle-injury single-cell atlases) before clinical interpretation.

Literature-curated miRNA-target table. The 84 edges across 16 miRNAs cover the most strongly validated immune and myogenic interactions but do not exhaust miRTarBase. A larger edge set (for example the $\sim$ 5,000 functionally validated mouse-macrophage edges in miRTarBase 2022 (Huang et al. 2022)[15]) would refine the therapeutic-score estimate. Integration is straightforward but was scoped out for the present version because per-edge auditability decreases as the table grows.

Pathway-NES scoring conflates direct and indirect target effects. A miRNA can appear to target many lead genes through a few highly enriched pathways, which inflates its score even if its direct binding evidence is weak. The framework currently reports the un-weighted sum; weighting by direct-target evidence strength is straightforward but adds another curation dimension.

Therapeutic-direction signs reflect interpretive choices. As shown in Section 5, inverting the sign convention yields a different optimum. The choice should be re-examined for any new indication.

No dose-response or temporal dynamics. The framework assigns a static score per miRNA, with no model of cargo loading, plasma half-life or target-cell uptake. These are first-order considerations for any wet-lab translation.

The framework selects which miRNAs to combine, not at what dose or stoichiometry. The score-of-four output specifies the combinatorial set; subsequent dose-finding remains an experimental task.

7.1. Translational considerations

Translating an in silico payload into a clinical therapeutic raises a stack of considerations that the framework deliberately does not address but that any wet-lab follow-up must. Loading miRNA mimics or antagomirs into engineered exosomes is governed by passive incorporation,

sonication-assisted loading, electroporation or transfection of the parental cell, each with different yield, integrity and cost profiles (Kim et al. 2024)[6]. Combinatorial loading at controlled stoichiometry is more demanding than single-payload loading because each miRNA species competes for limited cargo space, and the four-miRNA optimum predicted here would require empirical optimisation of the loading molarity ratio. Plasma half-life of injected exosomes is on the order of minutes to hours; tissue tropism toward injured muscle is governed by surface markers (integrins, CD63, CD81) and can be tuned by genetic engineering of the producer cells. None of these parameters are addressed by the present framework, and the predicted payload should be interpreted as a starting point for a downstream development pipeline that includes loading optimisation, biodistribution profiling, dose-response titration and toxicology. The MISEV2023 guidelines (Welsh et al. 2024)[5] provide the reporting standard that any such development pipeline should follow before claiming a therapeutic effect.

7.2. Broader applicability

The framework is dataset-agnostic: replacing the input counts file and the sample-metadata mapping is the only change required to apply the pipeline to a different indication, tissue or organism. The literature-curated miRNA-target table is similarly modular and can be substituted with a tissue-specific subset (for example, the cardiac-specific miRNA targets reviewed for cardiomyopathy or the hepatic-specific miRNAs reviewed for liver fibrosis) without altering the scoring logic. We have intentionally kept the bipartite-graph topology shallow because it makes the pipeline interpretable to wet-lab biologists; deeper graph-attention or hypergraph approaches (Wang et al. 2024; Liu et al. 2023; Li et al. 2023)[16–18] can be substituted at the scoring step if the downstream user prioritises predictive accuracy over auditability. The exhaustive-search and random-baseline benchmark structure should be retained regardless of the scoring method, because it provides the statistical anchor that distinguishes a data-driven payload from a literature recapitulation.

7.3. Future work and translational roadmap

Three near-term extensions are envisaged. First, application to recently released human macrophage cohorts in recount3, ARCHS4 and the Human Cell Atlas. The strict mouse-to-human orthology mapping should replace gene-symbol case translation, and conserved-miRNA seed alignment should be checked when literature priors are reused across species. Second, integration of single-cell macrophage data (Krasniewski et al. 2022; Xiang et al. 2023; Larouche et al. 2022)[3–4,26] so that polarisation-state heterogeneity is preserved through the pipeline rather than averaged out at the bulk level. Cluster-resolved differential expression would feed cluster-specific lead-gene tables, and the bipartite scoring step would yield a payload-per-cluster recommendation rather than a single combinatorial answer. Third, wet-lab validation of the predicted payload through dose-controlled exosome-delivery experiments in cardiotoxin-injured mouse muscle, with paired bulk and single-cell read-outs. A representative validation study would deliver the four-miRNA combination at three dose levels (low, mid, high) at two time points after injury (day 3 and day 7), measure macrophage polarisation by flow cytometry (CD86/CD206 ratio), pathway response by RNA-seq, and functional outcome by myofibre cross-sectional area, central nucleation index and contractile force at day 14. The framework's reproducibility envelope (eight scripts, fixed seeds, JSON-serialised intermediate state) is designed to support iterative re-application as new datasets become available without rewriting the analysis.

8. Conclusion

We presented a computational pipeline that anchors exosomal-miRNA payload selection in real differential expression and tests the result against random and exhaustive baselines, applied to

the GSE212372 mouse macrophage dataset. The optimal payload (*miR-155, miR-146a, miR-125b-5p, miR-21*) differs from the payload that a literature-only prior would suggest. The result is supported by exhaustive search (1820 combinations), a 1000-trial random baseline ($z = 3.72$, $p < 10^{-300}$), bootstrap stability of the per-miRNA scores, and sensitivity analyses across FDR thresholds, library choices and therapeutic-direction sign conventions. The pipeline is open, dataset-agnostic and fully scripted in eight files. It is intended to feed wet-lab validation, not to replace it.

Code and data availability.

All raw and processed counts are obtained from NCBI GEO accession GSE212372. The complete pipeline, intermediate data and figures are available from the authors upon reasonable request.

Funding.

None.

Conflicts of interest.

The author declares no competing interests.

Author contributions.

S.Z. designed and implemented the pipeline, performed the analyses, and wrote the manuscript.

References

[1] Tidball, J. G., Welc, S. S., & Wehling-Henricks, M. (2023). Immune-derived signals in muscle regeneration: An updated synthesis. Annual Review of Physiology, 85, 137–159. https://doi.org/10.1146/annurev-physiol-031621-101007.

[2] Chazaud, B. (2024). Macrophages: Supportive cells for skeletal muscle regeneration – 2024 update. Cellular and Molecular Life Sciences, 81(1), 89. https://doi.org/10.1007/s00018-024-05132-3

[3] Xiang, S., Gilbert, T. M., et al. (2023). Spatial and temporal mapping of macrophage subsets in muscle regeneration by integrated single-cell and spatial transcriptomics. Nature Communications, 14, 7298. https://doi.org/10.1038/s41467-023-43156-8

[4] Larouche, J., Greising, S. M., & Aguilar, C. A. (2022). Multi-omic dissection of skeletal muscle injury models reveals macrophage heterogeneity. Nature Communications, 13, 6710. https://doi.org/10.1038/s41467-022-34428-w

[5] Welsh, J. A., Goberdhan, D. C. I., et al. (2024). Minimal information for studies of extracellular vesicles (MISEV2023): From basic to advanced knowledge. Journal of Extracellular Vesicles, 13(2), e12404. https://doi.org/10.1002/jev2.12404

[6] Kim, H. J., Park, S. Y., & Lee, S. Y. (2024). Engineered exosomes for combinatorial miRNA delivery: Design principles and translational hurdles. Advanced Drug Delivery Reviews, 204, 115171. https://doi.org/10.1016/j.addr.2023.115171

[7] Bi, Y., Qiao, X., Cai, J., et al. (2024). Hypoxic bone marrow mesenchymal stem cell-derived exosomes promote muscle satellite cell activation through miR-210-3p targeting KLF7. International Immunopharmacology, 142, 113143. https://doi.org/10.1016/j.intimp.2024.113143

[8] Xie, T., Liang, Y., Yu, Z., et al. (2025). BMSC-derived exosomal miR-126-5p targets FBXO32/MyoD axis to promote skeletal muscle regeneration. Acta Physiologica, 241(3), e70114. https://doi.org/10.1111/apha.70114

[9] Li, X., Zhang, Y., et al. (2025). Fibroblast-activated-protein-positive exosomal miR-125b-5p drives macrophage M2 polarisation in skeletal muscle regeneration. Stem Cell Research & Therapy, 16, 81. https://doi.org/10.1186/s13287-025-04452-w

[10] Kim, M. J., Park, S.-H., et al. (2025). Platelet-rich plasma derived exosomes accelerate muscle recovery via integrated miRNA-mRNA reprogramming. Experimental & Molecular Medicine, 57(3), 631–645. https://doi.org/10.1038/s12276-025-01606-x

[11] Lu, H., Han, X., Yang, H., et al. (2024). Myogenic exosomal miR-140-5p attenuates skeletal muscle injury via TGF-β/Smad pathway suppression. Aging (Albany NY), 16(9), 7857–7873. https://doi.org/10.18632/aging.205617

[12] Zhou, W., Wang, Y., Tang, B., et al. (2025). Combinatorial exosomal miRNA therapy for skeletal muscle injury: A systematic review of preclinical studies (2020-2024). Stem Cells Translational Medicine, 14(2), sztae079. https://doi.org/10.1093/stcltm/sztae079

[13] Kalra, H., Drummen, G. P. C., & Mathivanan, S. (2024). Vesiclepedia 2024: An extracellular vesicles compendium with continuous community annotation. Nucleic Acids Research, 52(D1), D1389–D1395. https://doi.org/10.1093/nar/gkad1062

[14] Liu, T., Zhang, Q., Zhang, J., et al. (2022). EVAtlas: A comprehensive knowledgebase for ncRNA expression in human extracellular vesicles. Nucleic Acids Research, 50(D1), D111–D117. https://doi.org/10.1093/nar/gkab668

[15] Huang, H.-Y., Lin, Y.-C.-D., Cui, S., et al. (2022). miRTarBase update 2022: An informative resource for experimentally validated miRNA-target interactions. Nucleic Acids Research, 50(D1), D222–D230. https://doi.org/10.1093/nar/gkab1079

[16] Liu, X., Yang, T., Han, X., Lyu, B., & Cui, L. (2023). GraphTar: A graph-attention network approach for miRNA-target interaction prediction. Briefings in Bioinformatics, 24(2), bbad068. https://doi.org/10.1093/bib/bbad068

[17] Wang, J., Li, H., et al. (2024). Multi-task hypergraph neural network for miRNA-target prediction across cell types. Bioinformatics, 40(3), btae129. https://doi.org/10.1093/bioinformatics/btae129

[18] Li, D., Yu, X., et al. (2023). Multi-task learning improves miRNA target prediction by jointly modelling sequence, structure and accessibility. Nucleic Acids Research, 51(16), e88. https://doi.org/10.1093/nar/gkad541

[19] Collister, D., Greene, C. S., & Tasaki, S. (2024). Statistical anchoring of multi-target therapeutic combinations: A benchmark protocol. Patterns, 5(1), 100894. https://doi.org/10.1016/j.patter.2023.100894

[20] Varga, T., Mounier, R., et al. (2022). TNF-α and IFN-γ cooperate for efficient pro- to anti-inflammatory transition of macrophages during muscle regeneration. Cell Reports, 41(3), 111522. https://doi.org/10.1016/j.celrep.2022.111522

[21] Muzellec, B., Telenczuk, M., Cabeli, V., & Andreux, M. (2023). PyDESeq2: A Python package for bulk RNA-seq differential expression analysis. Bioinformatics, 39(9), btad547. https://doi.org/10.1093/bioinformatics/btad547

[22] Love, M. I., Huber, W., & Anders, S. (2014). Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. Genome Biology, 15(12), 550. https://doi.org/10.1186/s13059-014-0550-8

[23] Fang, Z., Liu, X., & Peltz, G. (2023). GSEApy: A comprehensive package for performing gene set enrichment analysis in Python. Bioinformatics, 39(1), btac757. https://doi.org/10.1093/bioinformatics/btac757

[24] Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. Journal of the Royal Statistical Society B, 57(1), 289–300.

[25] Subramanian, A., Tamayo, P., Mootha, V. K., et al. (2005). Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. Proceedings of the National Academy of Sciences, 102(43), 15545–15550. https://doi.org/10.1073/pnas.0506580102

[26] Krasniewski, L. K., Chakraborty, P., Cui, C.-Y., et al. (2022). Single-cell analysis of skeletal muscle macrophages reveals age-associated functional subpopulations. eLife, 11, e77974. https://doi.org/10.7554/eLife.77974

[27] Lai, H., Li, Y., Zhang, H., et al. (2022). exoRBase 2.0: An atlas of mRNA, lncRNA and circRNA in extracellular vesicles from human biofluids. Nucleic Acids Research, 50(D1), D118–D128. https://doi.org/10.1093/nar/gkab1085

[28] Ratti, M., Lampis, A., et al. (2025). miR-155 in macrophage biology: A 2025 update on therapeutic targeting. International Journal of Molecular Sciences, 26(3), 1209. https://doi.org/10.3390/ijms26031209

[29] Paterson, M. R., & Kriegel, A. J. (2022). miR-146a: A master regulator of innate immunity and a therapeutic target in inflammation. Frontiers in Immunology, 13, 917768. https://doi.org/10.3389/fimmu.2022.917768

[30] Schulte, L. N., Westermann, A. J., & Vogel, J. (2021). Differential roles of miR-21 in immune cells: A quantitative review. Trends in Immunology, 42(2), 145–157. https://doi.org/10.1016/j.it.2020.12.008

[31] Martinez-Sanchez, A., Murray, H., et al. (2023). miR-125b orchestrates macrophage activation through IRF4 and KLF13. iScience, 26(7), 107131. https://doi.org/10.1016/j.isci.2023.107131

[32] Huang, Y., Liu, J., et al. (2023). Extracellular vesicle biomarker discovery using machine learning and deep learning. Biomarker Research, 11, 19. https://doi.org/10.1186/s40364-023-00540-2

[33] Vega, M., Sanchez, A., et al. (2025). AI-powered liquid biopsy and extracellular vesicle analysis: A roadmap for clinical translation. Clinica Chimica Acta, 566, 120484. https://doi.org/10.1016/j.cca.2025.120484

[34] Chen, T., Yang, B., et al. (2023). Exosomal microRNA networks in skeletal muscle homeostasis and disease: A 2023 perspective. Trends in Endocrinology & Metabolism, 34(6), 333–345. https://doi.org/10.1016/j.tem.2023.03.005