

CROSS-SESSION BCI TRANSFER:

A MATCHED COMPARISON OF GLOBAL AND SELECTIVE POOLING

Yiming Shen David Degras

Department of Mathematics, University of Massachusetts Boston

Global Source Pooling (GSP) is the accuracy baseline | **Minimum Distance Pooling (MDP)** helps under source heterogeneity | **Bridge Domain Pooling (BDP)** supports robust domain adaptation under severe shift

Key takeaway: On BNCI2014_004, GSP gives the strongest matched final accuracy, while **MDP** and **BDP** become useful under source heterogeneity.

Background & Key Terms

- A BCI decodes user intention (e.g., motor imagery) from EEG.
- Each recording **session** can look different because EEG is non-stationary. [3]
- In deployment, we usually have many **labeled historical sessions** but only a short **unlabeled target session**. [2]
- The practical goal is to transfer past data safely to the new session.

Session Block	One recording block
Domain Shift	Source and target distributions differ
Selective Pooling	Retain target-relevant source sessions
DA Lift	$Acc_{DA} - Acc_{base}$ (secondary metric)
Matched Block	Same subject/method, different pipeline
Tuning Params λ	DA and classifier hyperparameters

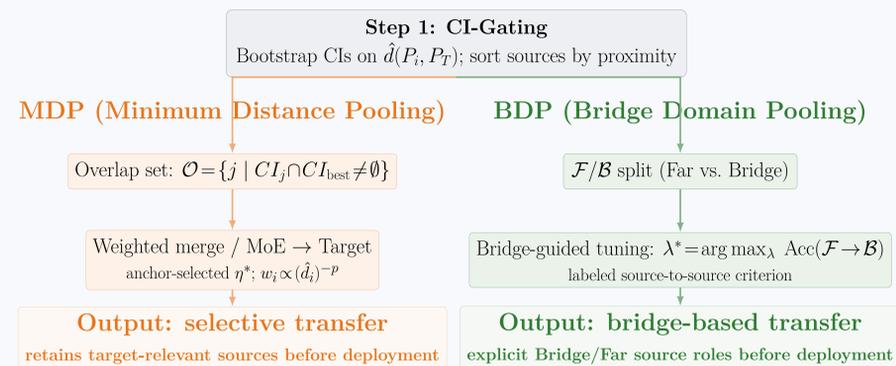
Covariance-based features (LogVar/TS) are standard in Riemannian BCI pipelines. [1,5]

Abbreviations (Quick Reference)

CI	Confidence interval (uncertainty in distance estimate)
CI-Gating	Select sources whose distance CIs overlap the closest source
CSP	Common Spatial Patterns (supervised spatial filtering)
GSP	Global Source Pooling (accuracy baseline)
MDP	Minimum Distance Pooling (Distance Gated Source)
BDP	Bridge Domain Pooling (robust DA under severe shift)

Method Overview

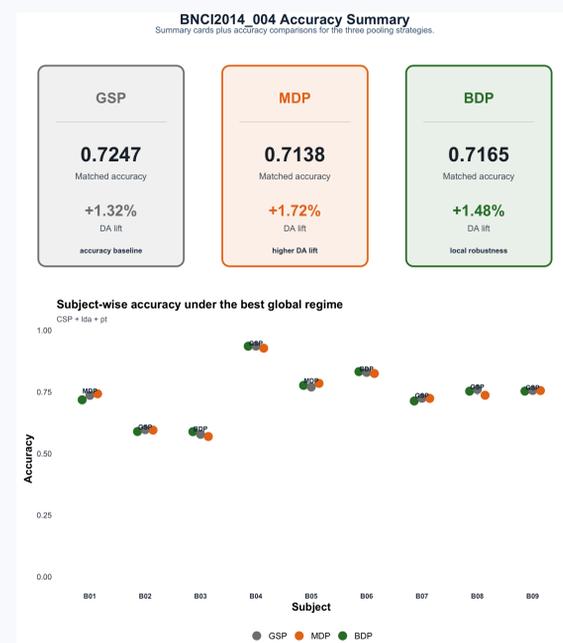
Intuition: Global Source Pooling (GSP) merges all historical sessions. **Minimum Distance Pooling (MDP)** retains source sessions that appear close to the target. **Bridge Domain Pooling (BDP)** assigns explicit Bridge and Far roles before transfer and uses Bridge-guided tuning for robust, conservative DA under severe shift.



Decision rule: (1) CI-gating yields Bridge size / uncertainty neighborhood. (2) Use MDP for Minimum Distance selection; use BDP when explicit Bridge/Far source roles are desired.

Key Results

Primary evidence: matched final accuracy on BNCI2014_004. DA lift is reported as a secondary within-strategy metric.



Interpretation of DA lift: $Acc_{DA} - Acc_{base}$ measures the benefit of turning on adaptation within the same strategy. It is **not** the primary basis for cross-strategy superiority claims.

Best configuration by strategy:

- Global Source Pooling (GSP): CSP + Linear SVM, DA=SA → Acc **0.7450**, Lift +3.34%
- Minimum Distance Pooling (MDP): CSP + Elastic Net, DA=SA, Dist=MMD → Acc **0.7416**, Lift +3.91%
- Bridge Domain Pooling (BDP): CSP + Elastic Net, DA=SA, Dist=MMD → Acc 0.7364, Lift +4.90%

Where selective pooling beats global pooling

Accuracy: only under Log-Variance + LDA + PT and Log-Variance + Linear SVM + PT.

Secondary DA lift: BDP reaches the largest single lift, +4.90% under CSP + Elastic Net + SA; MDP beats GSP in 12/24 matched methods on lift.

Key Comparisons

- Global Source Pooling (GSP) gives the highest matched final accuracy.
- Selective pooling beats global pooling in accuracy only in two Log-Variance + PT regimes.
- MDP shows higher DA lift than GSP (+1.72% vs +1.33%; secondary endpoint).
- BDP provides local robustness, with accuracy gains in 2/24 matched regimes.
- DA choice explains more controllable variance than pooling identity ($\eta^2 = 2.05\%$ vs 0.21% ; secondary omnibus view).

On BNCI2014_004, where cross-session drift appears relatively mild, Global Source Pooling is the strongest accuracy baseline, while Minimum Distance Pooling and Bridge Domain Pooling are better motivated for more heterogeneous or more severely shifted targets.

Domain Adaptation

EEG drifts between sessions: $P_{source} \neq P_{target}$.

- Sources:** abundant but heterogeneous
- Target:** short, **unlabeled**
- Pool all ⇒ **negative transfer**. [4]

Why transfer tuning is different

DA tuning is target-dependent, but the target session is unlabeled. All pooling strategies therefore tune hyperparameters using labeled source-side criteria rather than on the true deployment target.

Pooling Strategies

Global Source Pooling GSP	Minimum Distance Pooling MDP	Bridge Domain Pooling BDP
What it does Merge all labeled sessions	What it does Retain source sessions near the target	What it does Split sources into Bridge and Far sets before transfer
What it wins Final matched accuracy	Primary strength Higher within-pipeline DA lift in selected regimes	Primary strength Explicit source roles and local upside
Evidence Mean accuracy 0.7248 Wins vs BDP: 22/24 Wins vs MDP: 22/24	Evidence Mean lift + 1.72% vs GSP: + 0.39% BH-adjusted $p = 0.0129$	Evidence Best single lift + 4.90% under CSP + Elastic Net + SA Accuracy wins vs GSP: 2/24
Role Use as the empirical accuracy baseline	Role Useful when only a subset of historical sessions is target-relevant and far sources may cause negative transfer	Role Risk-control for severe target outliers and drift; favors generalizable transfer over local overfitting

Strategy	Main Idea	Primary Strength	Supported Claim
Global Source Pooling	Merge all sessions	Final accuracy	Strongest matched baseline on BNCI2014_004
Minimum Distance Pooling	Retain target-relevant sessions	Secondary DA lift	Helpful in narrower heterogeneous settings
Bridge Domain Pooling	Split into Bridge/Far subsets	Explicit source roles	Provides a Bridge/Far structure for selective pooling

This panel is accuracy-first: Global Source Pooling wins the primary endpoint, while Minimum Distance Source Pooling is useful in narrower settings and BDP mainly changes how source roles are organized.

- [1] Barachant et al. Multiclass BCI classification by Riemannian geometry. *IEEE TBME*, 2011.
 [2] Jayaram et al. Transfer learning in brain-computer interfaces. *IEEE CIM*, 2016.
 [3] Shenoy et al. Towards adaptive classification for BCI. *J Neural Eng*, 2006.
 [4] Wang et al. Characterizing and avoiding negative transfer. *CVPR*, 2019.
 [5] Zaïmi et al. Transfer learning: A Riemannian geometry framework with applications to BCIs. *IEEE TBME*, 2017.

Variance Structure

Full variance decomposition shown; the primary endpoint remains matched final accuracy.

Metric	Subject	DA	Pooling	Feature	Classifier
Raw η^2	93.88%	1.95%	0.18%	0.18%	0.01%
F-statistic	1886.80	104.65	14.81	14.59	0.87

Subject dominates overall variations

Among controllable factors only, **DA choice** explains about **10.8x** more variance than pooling identity
 Based on raw η^2 : $\eta_{DA}^2 = 1.95\%$ vs. $\eta_{pooling}^2 = 0.18\%$