BULLETIN N° 235
ACADÉMIE EUROPEENNE
INTERDISCIPLINAIRE
DES SCIENCES

INTERDISCIPLINARY EUROPEAN ACADEMY OF SCIENCES

Lundi 6 mai 2019 :
à 16h
à l'Institut Henri Poincaré salle 314
11, rue Pierre et Marie Curie 75005 PARIS/Métro : RER Luxembourg

---

Conférence :
" Biologie quantitative de la communication entre neurones:
instabilité moléculaire et mémoire, du normal au pathologique"
par Antoine TRILLER,

membre de l'Académie des Sciences
Directeur de Recherche (CE) INSERM,
Institut de Biologie de l'École Normale Supérieure
(ENS . CNRS UMR8197 . Inserm U1024)

Notre Prochaine séance aura lieu le lundi 3 juin 2019 à 16h30
à l'Institut Henri Poincaré salle 314

Elle aura pour thème

Conférence:
" *Enigmes posées par la mémoire et l'apprentissage"*
par Vincent HAKIM
Directeur de Recherche CNRS, Équipe biophysique et neuroscience théoriques
Département de Physique –École Normale Supérieure
24, rue Lhomond   75231 Paris Cedex 05

ACADÉMIE EUROPÉENNE INTERDISCIPLINAIRE DES SCIENCES
INTERDISCIPLINARY EUROPEAN ACADEMY OF SCIENCES

mai 2019

N°235

TABLE DES MATIERES

Prochaine séance : lundi 3 juin 2019

Conférence:
" Enigmes posées par la mémoire et l'apprentissage"
par Vincent HAKIM
Directeur de Recherche CNRS, Équipe biophysique et neuroscience théoriques
Département de Physique –École Normale Supérieure
24, rue Lhomond   75231 Paris Cedex 05

ACADEMIE EUROPEENNE INTERDISCIPLINAIRE DES SCIENCES

Fondation de la Maison des Sciences de l'Homme, Paris.

Séance du Lundi 6 mai 2019/IHP 16h

La séance est ouverte à 16h **sous la Présidence de Victor MASTRANGELO** et en la présence de nos Collègues Gilbert BELAUBRE, Jean BERBINAU, Jean-Louis BOBIN, Gilles COHEN-TANNOUDJI, Ernesto DI MAURO, Françoise DUTHEIL, Claude ELBAZ, Michel GONDRAN, Irène HERPE-LITWIN, Claude MAURY, Marie-Françoise PASSINI, Jacques PRINTZ, René PUMAIN, Jean SCHMETS, Jean-Pierre TREUIL .

Etaient excusés :François BEGON, Jean-Pierre BESSIS, Bruno BLONDEL, Michel CABANAC, Alain CARDON, Juan-Carlos CHACHQUES, Eric CHENIN, Alain CORDIER , Daniel COURGEAU, Sylvie DERENNE, Jean-Félix DURASTANTI, Vincent FLEURY, Robert FRANCK, Jean -Pierre FRANCOISE, Dominique LAMBERT, Pierre MARCHAIS, Anastassios METAXAS, Jacques NIO, Edith PERRIER, Pierre PESQUIES, Denise PUMAIN, Michel SPIRO, Alain STAHL, Mohand TAZEROUT, Jean-Paul TEYSSANDIER, Jean VERDETTI.

Etait présent notre collègue membre correspondant Benoît PRIEUR

I.   Conférence d'Antoine TRILLER

A. **Présentation du conférencier par notre Président Victor MASTRANGELO**.

Antoine TRILLER membre de l'Académie des Sciences, est Directeur de Recherche CE (DR) de l' U1024 - INSERM à l'Institut de Biologie de l'ENS (IBENS) .

Formation- Carrière

Après un doctorat en médecine en 1978, il a soutenu en 1985 une thèse d'état en Sciences Naturelles à l'UMPC. En 1994, il est devenu DR en en physiologie. En 1995, il a intégré en tant que DR INSERM, le département de Biologie de l'ENS. Puis de 2006 à 2014 il est devenu Directeur du Département de Biologie de l'ENS et de 2010 à 2018 Directeur de l'Institut de Biologie de l'ENS (IBENS). Son centre d'intérêt a toujours été le domaine de la biologie cellulaire des synapses et des conséquences de leurs dysfonctionnements dans les "synopathies" comprenant l'épilepsie, la douleur neuropathique et les maladies neurodégénératives.

Autres expériences et appartenances professionnelles

–   2007-2011   Président du SAB de "l'Institut de la Recherche sur la Moelle et l'Encephale
–   2008-2010   Co-Président équipe Scientifique de la "Fondation pour la Recherche Médicale"
–   2009-2010   Président Comité Atipe-Avenir (CNRS-Inserm) pour les neurosciences
–   2009-2011   Président du comité pilote de l'ANR pour les Maladies Neurologiques et Psychiatriques
–   2011-2013   Président du Conseil Scientifique de la "Fondation pour la Recherche Médicale" **jusqu'en 2015:**
–   Doyen de Recherche de l'"Université de Recherche PSL"

- Membre du groupe "Wellcome Trust Molecular Neuroscience Review "
- Vice-Président du fond de subvention "ERC LS5 Neuroscience Starting "

Actuellement :

- Directeur de l'"Institut de Biologie" de l'ENS " (IBENS)-unité Inserm-CNRS depuis 2010 Directeur des Labex "Memolife" associant IBENS, CIRB (Collège de France) et ESPCI
- Membre du Comité de Recherche de la "Fondation pour la Recherche Médicale"
- Membre du Comité Scientifique Consultatif du MPI pour la Recherche sur le Cerveau (Franckfurt)
- Membre du Comité Scientifique Consultatif du ZMBH, (Université de Hambourg)
- Membre du Jury pour le Prix Kavli en Neurosciences (Norvège) Membre du Jury du Prix Jeantet en Médicine
- Membre du Comité de Direction de la Fondation Bettencourt 'Agir pour l'Audition'
- Membre du Comité de Direction de l'Institut Pasteur

Distinctions:

- 2004 Grand Prix de l'Académie des Sciences: Prix du CEA
- 2004 Prix Inserm pour "Physiologie and Physiopathologie"
- 2007 Professeur Honoraire de l'Université Nuremberg-Erlangen
- 2008 1ère Conférence au département de Pharmacologie David Smith,. Oxford
- 2009 Chevalier de l'ordre national des Palmes Académiques
- 2010 Grand Prix de l'Academie des Sciences: Prix Lamonica en Neurologie
- 2011 Membre de l'Academie des Sciences
- 2011 Prix pour la Recherche Médicale de l'Association France Alzheimer
- 2013 Membre du EMBO
- 2015 Conférence Alfred Fessard
- 2017 Chevalier de la Légion d'Honneur

Thématiques principales de Recherche:

Il a fait plusieurs contributions autant pour comprendre les synapses que pour souligner les principaux traits uniquement associés aux synapses d'inhibition. Il a montré que les récepteurs du système nerveux central (CNS) s'accumulent aux synapses, en utilisant pour la première fois la microscopie électronique. Cette recherche a fourni un modèle pour la localisation par microscopie électronique autant du récepteur GABA ou acide γ-aminobutyrique que des récepteurs de glutamate. Avec Daniel Choquet, il a montré que les récepteurs de glycine (GlyRs) diffusent dans la membrane plasmatique , et que ceci est régulé par la géphyrine une protéine d'échafaudage cytoplasmique. Ceci a conduit vers un nouveau paradigme pour la synapse dans laquelle les composants membranaires forment un **complexe métastable** très dynamique (plutôt qu'un **complexe stable)** et diffusent de manière continuelle dans le plan de la membrane et en faisant du recyclage entre les compartiments internes et externes . Il a également développé des techniques d'analyse quantitative du comportement des biomolécules basées sur la microscopie de très haute résolution et le repérage des molécules individuelles **.** En 2003, il a mis en évidence le **premier repérage de molécule individuelle**, en utilisant des **récepteurs uniques de glycine marqués avec des points quantiques**, pour caractériser leur dynamique latérale dans la membrane neuronale des cellules vivantes . Plus récemment, en utilisant la microscopie de très haute résolution (PALM), il a développé une approche quantitative nanoscopique pour dénombrer les molécules et suivre leur dynamique dans les post-synapses individuelles .

Principales Publications:

1. Triller A, Cluzeaud F, Pfeiffer F, Betz H, Korn H. (1985) Distribution of glycine receptors at central synapses: an immunoelectron microscopy study. *Journal of Cell Biology*. 101: 683-8. PMID 2991304

2. Meier J, Vannier C, Serge A, Triller A, Choquet D (2001). Fast and reversible trapping of surface glycine receptors by gephyrin. *Nature Neuroscience*. 4: 253-60. PMID 11224541

3. Dahan M, Levi S, Luccardini C, Rostaing P, Riveau B, Triller A (2003) Diffusion dynamics of glycine receptors revealed by single-quantum dot tracking. *Science.* 302: 442-5. PMID 14564008

4. Specht CG, Izeddin I, Rodriguez PC, El Beheiry M, Rostaing P, Darzacq X, Dahan M, Triller A* (2013) Quantitative nanoscopy of inhibitory synapses: counting gephyrin molecules and receptor binding sites. *Neuron* 79: 308-321. PMID 23889935

5. Korn H, Triller A, Mallet A, Faber DS (1981) Fluctuating responses at a central synapse: n of binomial fit predicts number of stained presynaptic boutons. *Science*. 213: 898-901. PMID 6266015

6. Triller A, Cluzeaud F, Korn H. (1987) gamma-Aminobutyric acid-containing terminals can be apposed to glycine receptors at central synapses. *J. Cell Biol*. 104: 947-56. PMID 3031086

7. Levi S, Chesnoy-Marchais D, Sieghart W, Triller A (1999) Synaptic control of glycine and GABA(A) receptors and gephyrin expression in cultured motoneurons. *J. Neurosci*. 19: 7434-49. PMID 10460250

8. Racca C, Gardiol A, Triller A. (1997) Dendritic and postsynaptic localizations of glycine receptor alpha subunit mRNAs. The *J. Neurosci*.17: 1691-700. PMID 9030628

9. Gardiol A, Racca C, Triller A.(1999) Dendritic and postsynaptic protein synthetic machinery. *J.Neurosci*. 19: 168-79. PMID 9870948

10. Rosenberg M, Meier J, Triller A, Vannier C (2001). Dynamics of glycine receptor insertion in the neuronal plasma membrane. *J. Neurosci*. 21: 5036-44. PMID 11438579

11. Renner M, Wang L, Lévi S, Hennekinne L, Triller A. (2017) A Simple and Powerful Analysis of Lateral Subdiffusion Using Single Particle Tracking. *Biophysical J*. 113: 2452-2463. PMID 29211999

12. Bannai H, Lévi S, Schweizer C, Dahan M, Triller A. (2006) Imaging the lateral diffusion of membrane molecules with quantum dots.. Nature Protocol. 1(6):2628-34. PMID: 17406518.

13. Patrizio A, Renner M, Pizzarelli R, Triller A, Specht CG. (2017) Alpha subunit-dependent GlyR clustering and regulation of synaptic occupancy. *Scientific Reports*. 7: 10899. PMCID: PMC5589798

14. Lévi S, Schweizer C, Bannai H, Pascual O, Charrier C, Triller A (2008). Homeostatic regulation of synaptic GlyR numbers driven by lateral diffusion. *Neuron*. 59: 261-73. PMID 18667154

15. Bannai H, Lévi S, Schweizer C, Inoue T, Launey T, Racine V, Sibarita JB, Mikoshiba K, Triller A (2009). Activity-dependent tuning of inhibitory neurotransmission based on GABAAR diffusion dynamics. *Neuron*. 62: 670-82. PMID 19524526

16. Specht CG, Grünewald N, Pascual O, Rostgaard N, Schwarz G, Triller A (2011) Regulation of glycine receptor diffusion properties and gephyrin interactions by protein kinase C. *Embo J*. 30: 3842-53.PMID 21829170

17. Gouzer G, Specht CG, Allain L, Shinoe T, Triller A (2014). Benzodiazepine-dependent stabilization of GABA(A) receptors at synapses. *Mol Cell Neurosci*. 63: 101-13. PMID 25466558

II.   **Conférence** "*Biologie quantitative de la communication entre neurones: instabilité moléculaire et mémoire, du normal au pathologique.*"

Résumé de la conférence:

"Les neurones communiquent entre eux et forment des réseaux qui sont à l'origine des propriétés du système nerveux. Cette communication se fait au niveau de jonctions appelées "synapses". Les composants moléculaires des synapses sont instables et bougent tout le temps par des processus de diffusion. Ceci est à l'origine d'un paradoxe que nous aborderons: comment assurer une mémoire stable quand toutes les molécules bougent. Nous évoquerons ensuite les conséquences des altérations de ces mouvements dans des pathologies neurodégénératives comme la maladie d'Alzheimer ou de Parkinson."

Un compte-rendu rédigé par un membre de l'AEIS qui est neurobiologiste, sera prochainement disponible sur le site de l'AEIS http://www.science-inter.com.

Documents

− Pour préparer la conférence de Vincent HAKIM , nous vous proposons :

    p. 08 : un résumé de sa conférence *" Enigmes posées par la mémoire et l'apprentissage"*

    p. 09 : Un article de Vincent HAKIM , Antoine TRILLER et al. intitulé An aggregation-removal model for the formation and size determination of post-synaptic scaffold domains. paru dans V. PLoS Comput Biol. 2017 Apr 24;13(4):e1005516. doi: 10.1371/journal.pcbi.1005516. Site: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5421815/

    p. 27 : un article de Vincent HAKIM et Boris BARBOUR et al. intitulé Cerebellar learning using perturbations. paru dans Elife. 2018 Nov 12;7. pii: e31599. doi: 10.7554/eLife.31599 Site : https://www.ncbi.nlm.nih.gov/pubmed/30418871

Résumé Article Vincent HAKIM:

Enigmes posées par la mémoire et l'apprentissage

Les synapses sont des structures biologiques importantes qui permettent la transmission d'information entre neurones et elles sont aussi censées être le siège de l'apprentissage et de la mémoire. Cependant, la persistance d'une mémoire pendant des années face au renouvellement en quelques heures des composants synaptiques moléculaires, demeure mystérieuse. De même, pendant l'apprentissage d'une tâche complexe, le retour sur le succès ou l'erreur est le plus souvent global et peu informatif. Reste incompris comment les forces de multiples synapses peuvent s'ajuster de façon appropriée dans ces conditions - un problème dit d'affectation (ou « credit-assignment » en anglais) -. Après un rappel de quelques éléments de biophysique synaptique et de quelques découvertes expérimentales significatives, je discuterai de nos travaux récents en collaboration avec les équipes d'A. TRILLER et B. BARBOUR à l'IBENS, en vue de la résolution de ces énigmes.

RESEARCH ARTICLE

An aggregation-removal model for the formation and size determination of post-synaptic scaffold domains

Jonas Ranft[1]☯, Leandro G. Almeida[2]☯, Pamela C. Rodriguez[2], Antoine Triller[2]*, Vincent Hakim[1]*

1 Laboratoire Physique Statistique, CNRS, Ecole Normale Supérieure, PSL Research University, Université Pierre et Marie Curie, Paris, France, **2** IBENS, Ecole Normale Supérieure, PSL Research University, CNRS, INSERM, Paris, France

☯ These authors contributed equally to this work.
* antoine.triller@ens.fr (AT); vincent.hakim@ens.fr (VH)

Abstract

The formation and stability of synapses are key questions in neuroscience. Post-synaptic domains have been classically conceived as resulting from local insertion and turnover of proteins at the synapse. However, insertion is likely to occur outside the post-synaptic domains and advances in single-molecule imaging have shown that proteins diffuse in the plane of the membrane prior to their accumulation at synapses. We quantitatively investigated this scenario using computer simulations and mathematical analysis, taking for definiteness the specific case of inhibitory synapse components, i.e., the glycine receptor (GlyR) and the associated gephyrin scaffolding protein. The observed domain sizes of scaffold clusters can be explained by a dynamic balance between the aggregation of gephyrin proteins diffusing while bound to GlyR and their turnover at the neuron membrane. We also predict the existence of extrasynaptic clusters with a characteristic size distribution that significantly contribute to the size fluctuations of synaptic domains. New super-resolution data for gephyrin proteins established the existence of extrasynaptic clusters the sizes of which are consistent with the model predictions in a range of model parameters. At a general level, our results highlight aggregation with removal as a non-equilibrium phase separation which produces structures of tunable size.

Author summary

Synapses mediate information transmission between neurons and are the physical support of memory. It has been realized that synapses are dynamic biological structures. Neurotransmitter receptors diffuse in the neuron membrane and synaptic scaffolding proteins are constantly renewed. We propose a biophysical model that links these different measured quantities for inhibitory synapse components and show how they determine the size of postsynaptic domains. The model predicts that synaptic scaffolds also exist extrasynaptically and that they contribute to fluctuations of synaptic sizes. We

confirm by super-resolution microscopy the existence of extrasynaptic scaffolds. Their measured size distribution agrees with the model predictions for specific parameters. The model should be helpful to better understand the dynamics of synapses and their possible levels of regulation.

Introduction

Synapses mediate transmission of information between neurons and are generally thought to be, at least in part, the support of memory. However, when investigated at the molecular scale, synapses appear as dynamic assemblies, the constituents of which are exchanged on timescales ranging from tens of minutes down to seconds [1]. This raises fundamental questions about the way memory is maintained [2]. On the postsynaptic side, neurotransmitter receptors are mostly inserted in the neuron plasma membrane at non-synaptic loci [3–5]. Single-particle imaging and tracking techniques have shown that they subsequently diffuse in the plasma membrane in and out of synaptic domains (see [1] for a review and references therein). The postsynaptic density (PSD) contains scaffolding proteins, which provide binding sites for the receptors and transiently stabilize them at the PSD. The strength of a synapse is determined by the number of receptors at the PSD at a given moment. This number depends on the number of receptor binding sites provided by the scaffolding proteins [6] as well as on the affinity between receptors and scaffolding proteins [7]. Therefore, the size of the PSD, i.e., the number of scaffold protein binding sites, is a key determinant of synaptic strength.

We have focused our study on inhibitory synapses for which most components have been identified and functionally characterized. The gephyrin scaffolding protein is the central structural component of inhibitory synapses [8]. In most cases, scaffolding proteins, including gephyrin [9, 10], are renewed on timescales of minutes to hours in the synaptic domain (for a review see [11]). The basic oligomeric form of a full-length gephyrin is that of a trimer, mediated via strong interactions of the N-terminal G-domains [12–14]. In eukaryotic cells, gephyrin trimers eventually undergo further oligomerization [15], supposed to underlie post-synaptic clustering in neurons. Gephyrin is present in the cellular cytoplasm but it can also diffuse just below the plasma membrane when bound to receptors [16].

In the PSD, gephyrin proteins form homo-multimeric scaffolds just below the plasma membrane. Precise data have been gathered in the past years about several biophysical parameters governing the dynamics of receptors and scaffolding proteins, as reviewed in [1]. GlyR diffusion constants and concentrations inside and outside inhibitory synapses have been determined (see [17] and references therein). The size distributions of PSD and gephyrin scaffolds have been measured as well [6].

A model for the synapse based on these data and linking them is needed to address the questions of synapse formation, maintenance and dynamics from a precise biophysical viewpoint. Co-expression in fibroblasts of GlyR and gephyrin is sufficient to generate at the plasma membrane clusters of sizes similar to that of the PSD at inhibitory synapses [9]. This finding has provided a motivation to investigate theoretically how receptor diffusion, association of receptors to scaffolds, and the self-aggregation of scaffolding proteins could give rise to the formation of domains of given sizes. Some previous works have started to tackle this question. Reaction-diffusion equations for the above-described processes were proposed, and the authors suggested that a Turing-like instability could underlie the formation of synaptic domains [18, 19]. In another study, only receptors were described and it was alternatively proposed that synaptic receptor clusters could result from a dynamic balance between an

incoming flux of diffusing receptors and an efflux from the cluster to the cell cytoplasm mediated by receptor removal [20]. That diffusion, aggregation and removal of molecular components can serve to produce macromolecular domains in a membrane has been previously proposed for the formation and maintenance of lipid rafts and precisely examined in that context [21, 22]. Similar mechanisms have also been shown to control the organization or E-cadherin clusters at cell membranes [23]. It is clear that the lifetime of a cluster of proteins can be many orders of magnitude larger than the lifetime of its constituents, an attractive feature for a memory storing structure [2] which has been previously theoretically examined [24].

Based on these previous works, we have now examined the characteristics of scaffold domains produced by scaffolding proteins (here gephyrin) bound to receptors (here GlyR) diffusing in the plasma membrane. Taking into account the available biophysical data [6, 9, 16] we first deduce that the combined effect of aggregation, diffusion and removal produces typical scaffold domain sizes which are similar to those observed experimentally. Then, particle-based computer simulations and theoretical analyses allow us to comprehensively characterize the cluster sizes and dynamics resulting from these basic processes. In agreement with previous works on aggregation–removal in other contexts [25, 26], we show that aggregation and removal of scaffolding proteins provide a non-equilibrium process at the origin of the distribution of dynamically evolving scaffold domains of different sizes. Furthermore, we predict how the size distribution of these domains depends on biophysical parameters such as the turnover time of scaffolding proteins at the plasma membrane and the dependence of cluster diffusion with size.

The actual size distribution of gephyrin clusters was measured in cultured spinal cord neurons using super-resolution microscopy. By comparing these novel observations with model predictions, we then infer the biophysical parameters that govern the assembly of gephyrin domains according to the proposed model of protein aggregation, diffusion, and turnover.

Results

Balance between lateral aggregation and cytoplasmic recycling of scaffold proteins sets scaffold domain size

We set out to assess the combined role of

1. lateral diffusion of bound scaffold-receptor complexes in the membrane,

2. scaffolding protein aggregation, and

3. scaffolding protein removal at the membrane and in scaffolds

by first analyzing a reduced model of scaffolding protein dynamics at the membrane, see Fig 1. In this model, a single scaffold domain is surrounded by diffusing scaffolding "particles" (i.e. gephyrin trimers) bound to receptors. The domain edge acts as an absorbing boundary on the diffusing complexes, thus imposing a concentration gradient which in turn gives rise to an incoming diffusive flux of scaffold proteins.

Because this influx grows at most like the perimeter of the domain, while the protein efflux due to desorption of aggregated scaffolding proteins into the cytoplasm scales with the area of the domain, the balance of both fluxes occurs for a well-defined domain size. The resulting equilibrium size can be calculated analytically (see S1 Text: Single scaffold domain) and depends on biophysical parameters such as the diffusion coefficient of scaffold-receptor complexes in the extrasynaptic membrane D_0, the removal rate of scaffold proteins k, the total surface concentration of scaffold proteins at the membrane c_0, and the density of scaffold proteins within the domain ρ. Note that k is an effective rate that captures any local imbalance of
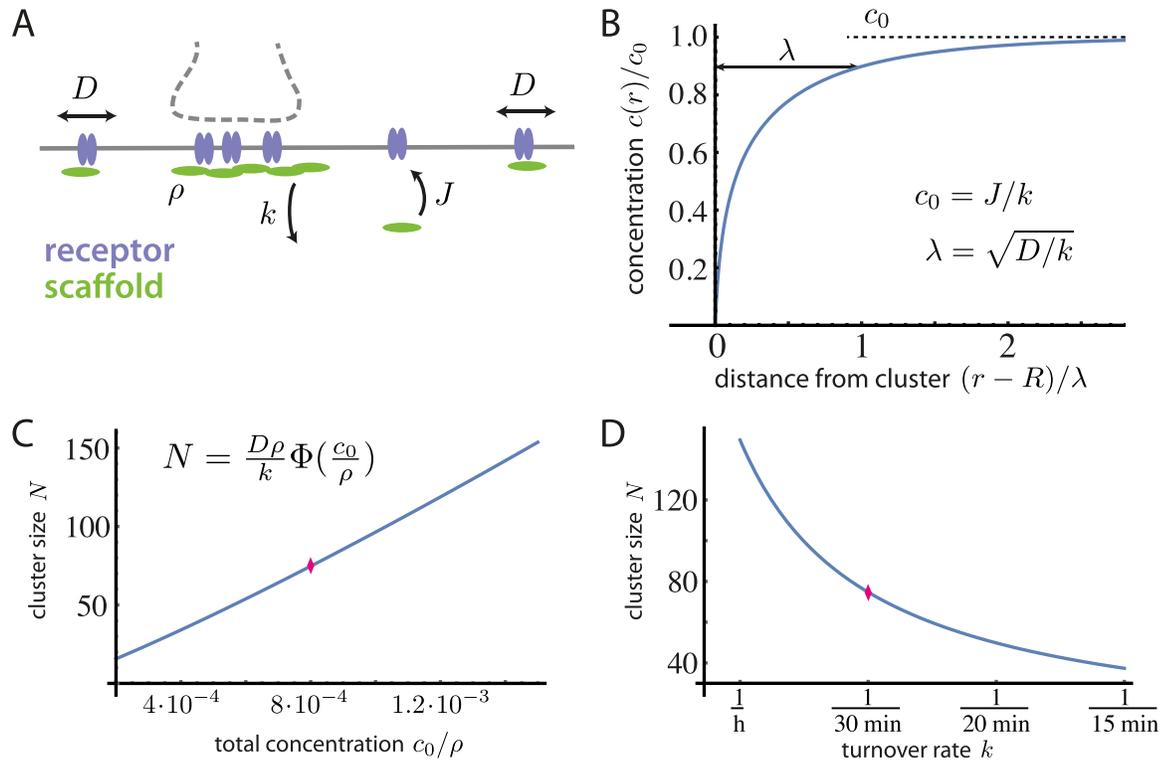
Fig 1. Lateral diffusion and aggregation of scaffold-receptor complexes at the post-synaptic membrane. A: Sketch of the model basic processes: receptors (GlyR, blue) and scaffold "particles" (gephyrin trimers, green) can form complexes; scaffold particles are brought to the cell membrane with a flux J and leave it with a rate k; sub-membraneous scaffold particles aggregate by homotypic scaffolding protein interactions; scaffold-receptor complexes diffuse laterally along the cell membrane, with a diffusion constant D. B: Concentration profile of diffusing scaffold particles around a disc-shaped scaffold domain of radius R. Far from the cluster, the concentration of diffusing scaffold particles is uniform and equal to J/k. The diffusing scaffold particles are depleted in a layer of size $\lambda = \sqrt{D/k}$ close to the cluster boundary. C: Dependence of the stationary domain size N on the particle concentration c_0. D: Dependence of N on the turnover rate k. Reference parameters are taken to be $D = 0.02 \ \mu m^2/s$ [16], $k = 1/(30 \ min)$ [9], $c_0 = 4/3 \ \mu m^{-2}$ [16, 27]. The concentration of scaffold particles/trimers in a post-synaptic domain is taken to be $\rho = 5000/3 \ \mu m^{-2}$ [6]. When parameters are varied in B and C, the reference values are marked (red losange).

binding and unbinding of scaffold particles from and into the cytoplasm, respectively. Measured in the number N of scaffold particles, its "building blocks", the expression for the domain size reads

$$N = \frac{\rho D_0}{k} \Phi\left(\frac{c_0}{\rho}\right). \tag{1}$$

(The explicit form of the function $\Phi(x)$ is given by Eq. (S5,S7) in S1 Text: Single scaffold domain.) In Fig 1C and 1D, the domain size N is shown as a function of c_0 and k, respectively.

To check the plausibility of scaffold formation and maintenance by diffusion and removal, we calculate the expected domain size using the above equation and parameter estimates for gephyrin and GlyR from the literature, see Fig 1. Assuming that the smallest occurring gephyrin unit or "building block" are gephyrin trimers [12–14], we obtain an approximate domain size of $N \simeq 70$ trimers, or 210 gephyrin monomers, which is surprisingly close to previously published measurements of gephyrin domain sizes [6]. We therefore conclude that diffusion, aggregation, and turnover may indeed be key mechanisms involved in setting the size of synaptic scaffold domains.

The above analysis explicitly relates the size of post-synaptic gephyrin domains to the measured diffusion constant of GlyRs and the turnover rate of gephyrin in scaffold domains. It is however based on the simple assumption of a single circular gephyrin domain surrounded by diffusing GlyR-gephyrin trimer complexes. To further investigate the proposed scenario of post-synaptic scaffold domain formation we proceed with studying a more detailed, truly particle-based model of scaffold aggregation by means of computer simulations. This allows us to test the influence of simplifying assumptions, like the assumed circularity of scaffold domains, and to address the possibility of multiple diffusing scaffold domains of various sizes. Most interestingly, this allows us to obtain predictions from the model on the size distribution of scaffold domains that we then compare to new data.

Particle-based model of scaffold-protein dynamics at the membrane

In the particle-based simulations, we consider individual scaffold particles attached to the membrane that diffusive laterally at the membrane. Particles bind to each other upon encounter during their random diffusive trajectories, mimicking the homotypic interactions of scaffold proteins.

Since the details of scaffold aggregation, diffusion, and domain dynamics are yet to be described, we choose to make simple assumptions to obtain a computationally efficient model with few parameters as detailed below (see also Materials and methods for the details of our implementation).

Particles aggregate upon encounter and form clusters. We assume that particles in a cluster rearrange into a circular disc-like domain shape. (To test the influence of this assumption, we also performed a few simulations where, on the contrary, no rearrangement was allowed, as described further below.)

The disc-like clusters can themselves diffuse and aggregate. We thus need to prescribe a possible size-dependence of the diffusion constant. The Saffman-Delbrück theory [28] for thermal diffusion would predict a logarithmic size dependence of the diffusion constant arising from hydrodynamic effects (see however [29] for proteins of size comparable to the membrane thickness). This classic result may be modified by possible interactions of the scaffolding proteins with the cell cortex, non-thermal effects as well as by the more complex nature of the receptor-mediated diffusion of scaffold domains. For simplicity and to avoid introducing a characteristic size "ad hoc", we consider a power-law size dependence

$$D(n) = D_0 n^{-\sigma}, \ \ \sigma \geq 0 \tag{2}$$

where $D(n)$ is the diffusion constant of clusters containing n scaffold particles and σ is the size-dependence exponent ($\sigma = 0$ when Saffman-Delbrück's result [28] applies). In our model, σ is a further parameter of the system in addition to those introduced previously.

Particles are supposed to desorb into the cytoplasm with an effective rate k, modeled by a stochastic removal of individual particles from the membrane. We assume a constant rate k, irrespective of the size of the domain that particles belong to. As shown below, this assumption appears sufficient to account for present data. We also neglected lateral desorption of particles in the membrane after aggregation since binding affinity between gephyrin trimers is high. This amounts to assuming that the concentration of diffusing gephyrin particles in the membrane is large as compared to the concentration in equilibrium with the condensed scaffold phase (i.e., we neglect gephyrin "vapor pressure"). The existence of a significant lateral desorption would effectively amount to a reduced incoming flux of particles on each cluster and in an underestimate of our present fitted diffusion constant. Lateral desorption

would also favor more compact aggregates [30] (see also Discussion) in the case of partial rearrangement dynamics.

The desorption of particles is balanced by an incoming flux J of single particles to the membrane, which ensures that the average total concentration of particles attached to the membrane $c_0 = J/k$ remains constant.

Extrasynaptic aggregation of synaptic proteins

A typical snapshot of our simulations is shown in Fig 2A. In going beyond the simplified assumption of a single scaffold non-diffusing domain, we notably find in our simulations the continuous generation of small clusters at the membrane. These clusters continue to diffuse, albeit more slowly because of their increased size, leading to an ensemble of clusters of all sizes (Fig 2A–2C). Eventually, the growth of a particular cluster is limited by the turnover of its constituent scaffold particles, analogous to the reduced model of a single domain discussed above.

As a rather general result, we find that clusters become progressively rarer with increasing size (Fig 2B and 2C). The size distribution of clusters for a diffusion constant exponent $\sigma = 0.5$ is shown in Fig 2B both when aggregating particles are fully rearranged into circular domains or, on the contrary, when no rearrangement is performed after aggregation. While rearrangement after clustering has a large effect on domain shape (see insert of Fig 2B), its impact on cluster size distribution is minor. This is found to be true also for other diffusion constant exponents. Since we do not focus here on scaffold domain shapes, for computational efficiency, we consider in the following only fully rearranged clusters.

The shape of the cluster size distribution is governed by the size-dependence of the cluster diffusion constant (Fig 2C). Qualitatively, when σ increases away from zero (i.e. the diffusion of large clusters is increasingly suppressed), the large size limit of the distribution decreases, implying a smaller range of cluster sizes. Clusters are also more evenly distributed among the available size range.

The cluster size distributions produced by particle simulations are quantitatively compared in Fig 2C to numerical solutions of Smoluchowski rate equations [25, 26, 31, 32] (see S1 Text: Rate Equations). These equations account quite accurately for the obtained distributions after fitting a single overall kinetic parameter (see S1 Text: Rate Equations and S1 Fig). For a size-independent diffusion constant, one can analytically show that the distribution is a power-law with an exponent of -3/2 and an exponential cut-off (Eq. S14 in S1 Text: Rate Equations). In the general case $\sigma \neq 0$, the rate equations still allow us to numerically determine the stationary solution for the cluster size distribution (see also S1 Fig). When $\sigma > 0$, diffusion is progressively slower for larger and larger clusters. This reduced diffusion limits the growth of large clusters relative to smaller clusters, and the distribution is shifted towards smaller sizes with increasing σ. The shape progressively deviates from the power-law observed for $\sigma = 0$, eventually leading to the appearance of a shoulder at a smaller cut-off size beyond which clusters become again exponentially rare.

Larger clusters are rarer than smaller ones, but since they contain more scaffold particles, the majority of scaffold particles may still be found in clusters of large size. A useful quantity for characterizing such distributions is the "typical" cluster size, which corresponds to the average cluster size when clusters are weighted by the number of particles they contain. It is defined as $\langle\langle n \rangle\rangle \equiv \sum_n n^2 c_n / \sum_n n c_n$, where c_n is the concentration of of clusters of size n.

The typical cluster size as a function of the biophysical parameters of diffusion and turnover is shown in Fig 2D. We find that for a given diffusion constant size-dependence exponent σ, the typical size essentially scales with the dimensionless parameter combination $c_0 D_0 / k$
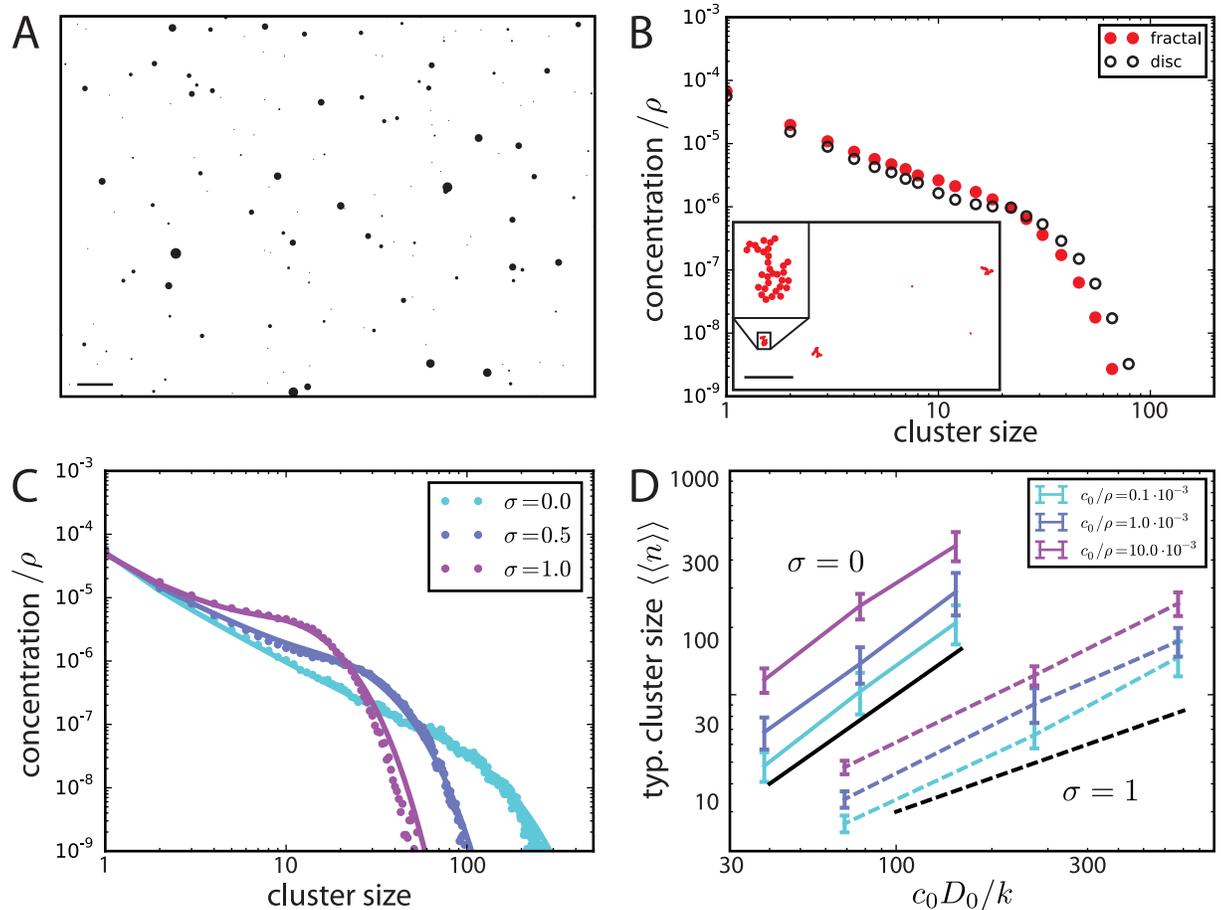
Fig 2. Simulations of scaffold domain formation. When not explicitly varied, simulation parameter values are $k = 2.0 \cdot 10^{-5} D_0 \rho$, $c_0 = 0.9 \cdot 10^{-3} \rho$, where $\rho = 0.77a^{-2}$ and a is the diameter of basic particles, and the diffusion exponent (Eq (2)) is $\sigma = 0.5$. A: Snapshot of scaffold domain dynamics. Scale bar 100 a; individual clusters are enlarged 3-fold for better visualization. B: The characteristic distribution of observed cluster sizes does not depend on the details of internal cluster structure. The cluster size distribution is shown for the simulation in panel A (black circles) with perfect particle rearrangement, i.e. disc-like clusters, and for a simulation without any particle rearrangement but otherwise identical parameters (red dots; simulation snapshot and typical cluster shown in inset). Note that the density of fractal clusters is not constant and depends on cluster size; for comparison the density $\rho = 0.77a^{-2}$ of disc-like clusters is used here. Scale bar inset 50 a. C: Characteristic cluster size distributions of simulations with different diffusion exponents σ but otherwise identical parameters (filled circles). The data for $\sigma = 0.5$ corresponds to the simulation shown in A. Superposed on the simulation results are the theoretical curves obtained from the rate equations (solid lines), see text for details. D: Scaling of the typical cluster size $\langle\langle n \rangle\rangle$ with $c_0 D_0/k$ for different σ and c_0. Solid black line: $\langle\langle n \rangle\rangle \propto c_0 D_0/k$; dashed black line: $\langle\langle n \rangle\rangle \propto (c_0 D_0/k)^{1/2}$.

https://doi.org/10.1371/journal.pcbi.1005516.g002

according to

$$\langle\langle n \rangle\rangle \propto \left(\frac{c_0 D_0}{k}\right)^{\alpha}, \quad \alpha \simeq \frac{1}{1+\sigma} . \tag{3}$$

For constant surface concentration of scaffold particles c_0 and a given σ, the typical cluster size increases with the general diffusion constant D_0 and decreases with higher effective turnover rate k. (Note that for constant cytoplasmic particle flux J onto the membrane, varying k also changes the total surface concentration via $c_0 = J/k$.) In the opposite case, for varying c_0 while keeping the product $c_0 D_0/k$ constant, the typical cluster size remains roughly constant; consistently, the scaling of $\langle\langle n \rangle\rangle$ with $c_0 D_0/k$ does not depend on the specific value of c_0 for which the simulations were done.

In principle, the scaling exponent α depends on the precise shape of the cluster size distribution, but a simple scaling argument rationalizes the observed relation (Eq 3). Assuming that the majority of scaffold particles are aggregated in clusters of size N, the concentration of these clusters is given by $c_N \simeq c_0/N$. The average distance between the clusters thus scales as $L \simeq \sqrt{N/c_0}$. The diffusion constant of the clusters being given by $D_N = D_0 N^{-\sigma}$, the timescale of their diffusive encounter scales as $T \propto L^2/D_N \simeq N^{1+\sigma}/(c_0 D_0)$. For N to be a stable cluster size, this timescale has to match the the typical turnover time $1/k$, otherwise clusters would either melt or grow bigger. One then immediately obtains $N \propto (c_0 D_0/k)^{1/(1+\sigma)}$.

Size distribution of gephyrin clusters

Our simulations predict that the sizes of scaffold domains depend on the continuous aggregation of scaffolding proteins into clusters at the membrane, which in turn supply the formation of larger domains. Because the expected distribution of cluster sizes is determined by the biophysical parameters of diffusion and turnover, these parameters can in principle be inferred from observed distributions of scaffold cluster sizes. Therefore, we experimentally determined the size distribution of gephyrin scaffolds in the neuronal membrane of cultured spinal cord neurons, both to confirm previous findings that extrasynaptic clusters exist and to provide a measurement of the relevant parameters D_0/k, c_0, and σ. To this aim, super-resolution microscopy on fixed neurons was used to count the number of individual gephyrin proteins in clusters (see Materials and methods, as previously published [6]). In brief, fluorescently labelled gephyrin proteins are stochastically activated at low light intensity such that the probability of concurrently activating close-by proteins is vanishingly small. The spatially separated light bursts of activated proteins can thus be localized with sub-wavelength precision. Because every labelled protein is activated one single time before being trapped in a "dark" state, the whole protein population can be imaged over the course of the experiment, and the clusters reconstructed from the recorded particle positions, see Fig 3A–3C.

The number of detections was directly associated to gephyrin cluster size [6]. We find gephyrin clusters of all sizes up to clusters of 71 trimers (see Fig 3E, S2 Fig). The largest domains are likely to be postsynaptic, 89% being found apposed to synaptic terminals in previous work [33]. Smaller gephyrin clusters should correspond to extrasynaptic clusters, consistent with the so-called nanoclusters previously visualized outside synapses with super-resolution microscopy [6].

The existence of clusters of different sizes was consistent with our predictions. Therefore, we compared the experimentally measured distributions to distributions generated by our model for different parameter combinations. Calculating the likelihood of the data for given theoretical distributions, we determined the parameters that explain best our data and which we refer to as maximum-likelihood fit of our model. Combining this analysis for all cultures ($n = 3$) by letting the total scaffold concentration, c_0, vary between cultures, we obtain a global estimate for the diffusion exponent, $\sigma = 0.5$, and the ratio of the diffusion constant over the particle turnover rate, $D_0/k = 30\ \mu\text{m}^2$, see Fig 3D. The total scaffold concentration varies among cultures from $c_0 = 0.5 - 5\ \mu\text{m}^{-2}$, see Fig 3E.

Cluster dynamics and size fluctuations of scaffold domains

Our simulations allow us to address the temporal size fluctuations of scaffold clusters and their stability over time. The size trajectory of an individual cluster is shown in Fig 4A. It fluctuates around a well-defined mean size, alternating between stochastic increases due to fusion with other clusters, and relatively smooth decreases on a characteristic timescale τ_{fluc} due to the continuous particle loss into the cytoplasm. This behavior is typical of all followed clusters. In
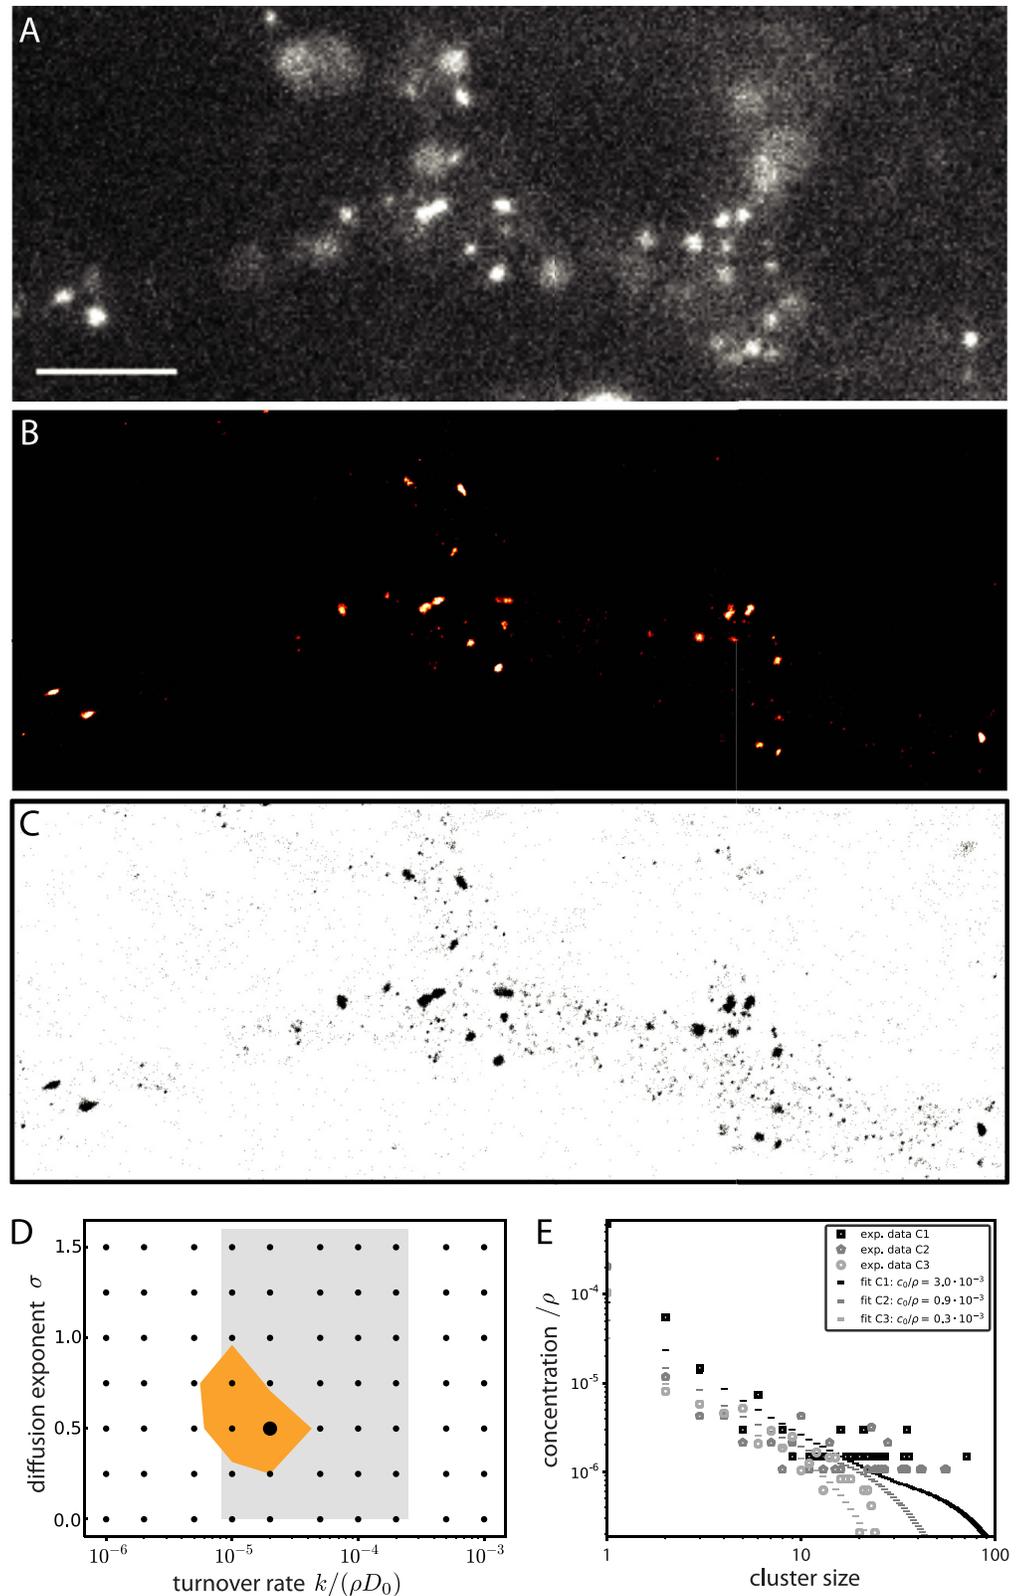
Fig 3. Gephyrin clusters and comparison with model distributions. A: Conventional fluorescence microscopy of spinal cord neurons expressing the mEos2-gephyrin protein. Scale bar 5 μm. B: Rendered super-resolution PALM image of the same segment shown in A acquired over 20k frames at a frame rate of

particular, clusters of small initial sizes grow by fusion with other clusters and fluctuate around the mean size after a few τ_{fluc} (Fig 4A, light blue lines). This leads the distribution of sizes explored by a cluster over time, shown in Fig 4B, to differ from the instantaneous distribution of all clusters at the membrane at a given time (Fig 3E), with a suppression of small cluster sizes. The cluster size autocorrelation $C_2(\tau) = \langle n(t + \tau)n(t) \rangle - \langle n \rangle^2$ allows a quantification of the fluctuation dynamics around the mean cluster size. It is shown in Fig 4C both for the particle-based simulations and for simulated trajectories based on a master equation approach
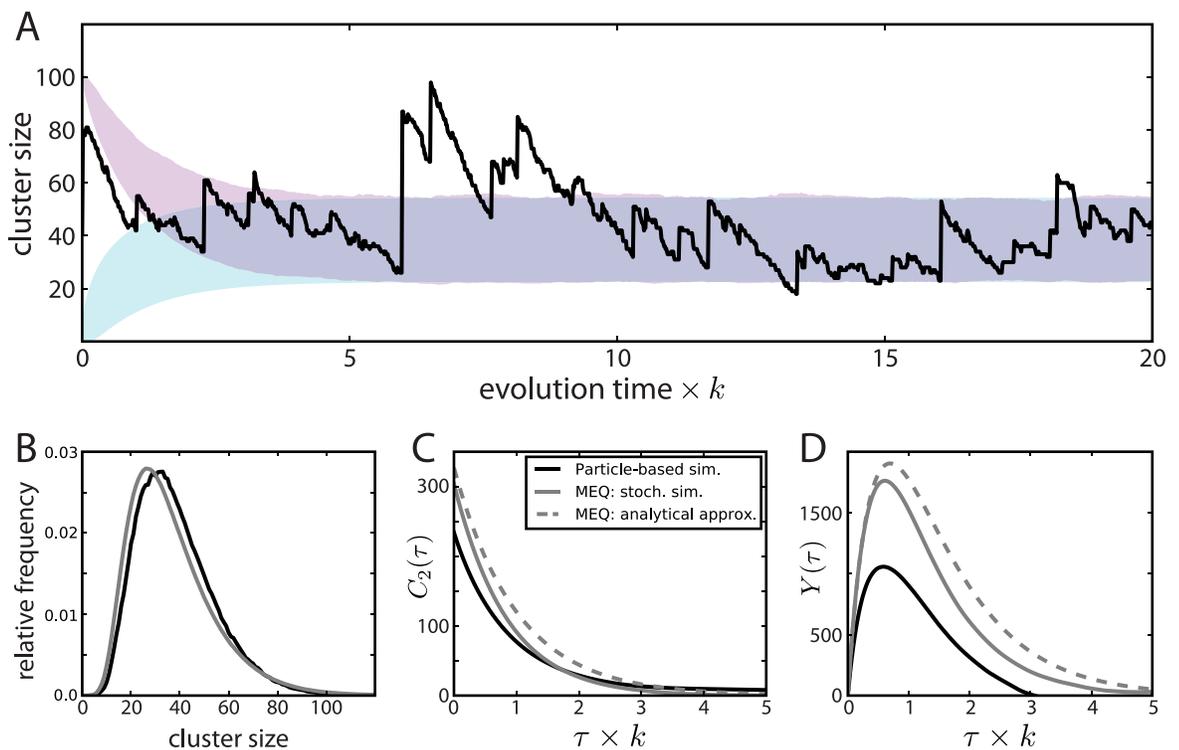
Fig 4. Predicted size evolution of individual clusters. A: Average cluster size (mean ± SD) evolution for clusters of an initial size of 2 (light blue) and 100 (light red) particles, respectively, averaged over independent simulated cluster size trajectories ($n_2 = 50 \cdot 10^6$, $n_{100} = 4000$) sampled from simulations. A typical cluster size trajectory is also shown (dark solid line). Simulation parameters are identical to Fig 2A. B: The distribution of cluster sizes at long times. Large clusters persist for a long time and explore the size distribution shown here due to particle addition by fusion with other clusters and particle removal by desorption. The distribution obtained from simulated trajectories ($n = 1000$) using the master equation (MEQ) approach (grey solid line, see text for details) is slightly different as spatial correlations are ignored. C: Autocorrelation, $C_2(\tau) = \langle n(t + \tau)n(t) \rangle - \langle n \rangle^2$, of the cluster size $n(t)$ obtained from full particle simulations (black solid line) and from the MEQ approach (solid grey line). The analytical prediction of an exponential decay with characteristic rate k is also shown (dashed grey line). D: The function $Y(\tau) = \langle n(t + \tau)n(t)^2 \rangle - \langle n(t + \tau)^2 n(t) \rangle$ is shown (solid black line). The fact that $Y(\tau)$ is non-zero implies that the dynamics is not invariant under time reversal and therefore out-of-thermodynamic-equilibrium. The approximation $Y(\tau) \propto \exp(-k\tau)[1 - \exp(-k\tau)]$ is also shown (dashed grey line) as well as the function $Y(\tau)$ generated by the MEQ approach (solid grey line) (see text and S1 Text: Rate Equations for details).

https://doi.org/10.1371/journal.pcbi.1005516.g004

(see S1 Text: Rate Equations). A simplified description of the latter approach provides an explicit expression for $C_2(\tau)$ and predicts that τ_{fluc} is set by the scaffolding protein turnover rate, $\tau_{\text{fluc}} \simeq 1/k$ (see S1 Text: Rate Equations) in good quantitative agreement with the simulations.

The cluster size variation over time depicted in Fig 4A is clearly not invariant under time reversal. It is interesting to note that it is a direct signature of the out-of-thermodynamic-equilibrium nature of the process considered since fluctuations at thermodynamic equilibrium should not allow to determine the arrow of time. One quantitative measure of the out-of-equilibrium nature of the underlying dynamics is provided by the third moment $Y(\tau) = \langle n(t + \tau)n(t)^2 \rangle - \langle n(t + \tau)^2 n(t) \rangle$ [34, 35] which by construction vanishes for systems that are invariant under time reversal. The non-trivial function $Y(\tau)$ obtained from cluster size time trajectories is shown in Fig 4D, together with approximations based on the master equation approach (see S1 Text: Rate Equations).

Discussion

We propose a model of scaffold domain formation and maintenance. It is based on

1. the delivery of scaffolding proteins at the plasma membrane,

2. receptor-mediated scaffold diffusion at the plasma membrane, and

3. scaffold growth upon encounter and multimerization with other scaffolds.

The counterbalance of this cluster growth process by the continuous removal of scaffolding proteins leads to a stabilization of a maximal domain size as well as a to a stationary distribution of scaffold clusters of different sizes. In a previous work [6], a population of small extrasynaptic clusters had been shown to coexist with large synaptic clusters, but the sizes of the former had not been precisely quantified. We have here provided new super-resolution microscopy data which shows that post-synaptic clusters assume a continuous range of sizes from gephyrin trimers up to sizes characteristic of synaptic domains. This agrees with a characteristic signature of the aggregation-removal dynamics proposed to underlie their formation. Comparison of experimental data with model size distributions allowed us to refine the biophysical parameters of this dynamics. Moreover, we found that the existence of diffusing extrasynaptic clusters has important consequences for the size fluctuations of synaptic domains, as the former may fuse with the latter. Our simulations have been focused on the dynamics of these extrasynaptic clusters. When one of them is followed over time, it reaches on the timescale of an hour the large size end of the domain size distribution and fluctuates around it, as shown in Fig 4A. In the simplest picture, the synaptic domains are simply the product of this formation process and their size distribution is as given by Fig 4B. However, this is most probably an oversimplification since the properties of synaptic scaffold domains certainly differ from extrasynaptic ones. For instance, we expect them to diffuse even less than large extrasynaptic scaffold domains. This would reduce their probability of encounters with other domains and consequently their size (simulations with one fixed domain show a mean size reduction of 33% as compared to Fig 4B).

In recent years, thermodynamic phase transitions and phase separation have been found to underlie the formation of various cellular membrane-less structures [36], most recently the PSDs of excitatory synapses [37]. One feature of thermodynamic phase separation is that the growth of the condensed phase is only limited by the depletion of the condensing component. The size of the formed structure is thus determined by the size of its "container" as well as the total number of molecular components enclosed in the latter [38]. In cells, the container can

be the whole cell or one of its sub-compartments, such as a spine for an excitatory PSD. Non-equilibrium phase separation may allow cells to avoid this constraint and to form localized structures of definite size, such as PSDs, irrespective of the size of the compartment in which they form, namely the whole neuron membrane for the majority of inhibitory synapses. Aggregation with removal or turnover has been proposed in various cellular contexts [21, 23] besides the one considered here. It is tempting to suggest that it is a specific, easily implementable, out-of-thermodynamic-equilibrium phase separation mechanism generally used by cells to form localized structures.

Our results expand previous works considering the non-equilibrium dynamics of receptors and scaffolds at the plasma membrane. While Haselwandter et al. [18, 19] considered general continuous equations for the interaction of receptors and scaffold proteins, they analyzed a different dynamical regime. Their mathematical analysis showed that domains of a characteristic size can be obtained via a Turing-like instability. In this case, the cooperative binding of scaffold proteins at synaptic domains is limited by receptors diffusing away from a domain and steric repulsion between receptors. However, it is unclear whether these assumptions are all verified in neurons, as for example, the receptor density appears to be far below saturation on synaptic domains [39, 40]. Moreover, in this Turing regime, a continuous range of domain sizes would not be expected. Burlakov et al. [20] proposed a model based on receptor aggregation and removal similar in spirit to the one described in Fig 1, but they did not consider scaffolding proteins nor cluster size distributions.

In the present study, we have studied a simple model of PSD formation, with a minimal number of parameters, most of which are determined by available data. Nonetheless, several of our assumptions may need to be refined when further data become available. In our model, receptors act mainly as carriers that enable the lateral diffusion of scaffold particles and domains along the cell membrane. Our single rate k for scaffold removal from the membrane accounts both for scaffold dissociation from receptors and for the endocytosis of receptor-scaffold complexes. This appears sufficient at present but receptors may need to be accounted for independently of scaffolds if one wishes to describe these two processes more precisely. A description with multiple molecular species would also be needed to describe the concentration of receptors on scaffold domains, its dependence on scaffold-receptor affinity as well as an influence of receptors of scaffold domain stability. A dependance of scaffold removal on scaffold domain size may also need to be included.

We have mainly considered the limit where scaffold domains rearranged themselves very efficiently and adopt a spherical compact shape. We have shown that this does not have significant consequences for cluster dynamics and size distributions by considering the other extreme limit of negligible cluster reshaping after scaffold aggregation, see Fig 2B. However, the shape of the individual scaffold clusters does depend on the rearrangement of scaffold particles within a domain or absence thereof (see e.g. inset Fig 2B). It is well-known that in the absence of recycling, clusters which grow by diffusion are prone to shape instabilities and that shapes are sensitive to changes in the dynamics. A prototypical example is provided by the classical diffusion-limited aggregation (DLA) model [41], where clusters grow by irreversible aggregation of very diluted single monomers without rearrangement. Such clusters assume isotropic ramified shapes with a fractal dimension $d_f \simeq 1.7$ for monomer diffusion in two dimensions ($d = 2$) as considered here. Preferred directions of attachment or crystalline anisotropy make the clusters anisotropic as commonly observed with snowflakes [42]. When not only monomers but also clusters diffuse and aggregate—a case known as cluster-cluster aggregation—the produced cluster shapes are different with a fractal dimension of $d_f \simeq 1.45$ [30]. Growth in a monomer solution of finite concentration c_0 renders DLA clusters compact above a characteristic length that depends on c_0 [30]. Removal of monomers furthermore introduces

a "diffusion length" $\lambda = \sqrt{D_0/k}$ above which clusters also cease to be fractal. With the parameters of Fig 3D, $\lambda \simeq 5\ \mu$m is much larger than the linear size of synaptic domains. It will therefore be an interesting task for future studies, both experimental and theoretical, to investigate scaffolding protein removal from scaffold domain and the modes of scaffold domain rearrangements and link them to the observed shapes which appear neither circular nor fully fractal.

The strength of a synapse is one of its key properties since it controls the efficacy of information transmission between neurons. However, when synapses are monitored over the course of days, they undergo large fluctuations in size and molecular content. This holds for neurons cultured *in vitro*, in absence of learning protocols, and even in absence of activity [43]. A phenomenological model, based on a mixed additive-multiplicative stochastic process [44], has been found to account well for the recorded size fluctuations of excitatory [45] as well as inhibitory [46] synapses. Interestingly, we found here that the existence of diffusing extrasynaptic clusters may provide a mechanistic explanation for similar fluctuations of scaffold domain sizes over shorter periods. The predicted distribution of fluctuating domain sizes is skewed (see Fig 4B) and qualitatively similar to those reported (see e.g. Fig 4 in ref. [46]).

Besides spontaneous fluctuations, the elucidation of the mechanisms by which synaptic strength is specifically modified remains a key challenge. The role of lateral diffusion of receptors in the plasma membrane in mediating receptor number changes at the synapse has started to be considered in the context of synaptic plasticity [1, 8] and synaptic strength homeostasis [47]. Both the lateral mobility of GlyR in the cell membrane and its binding affinity for gephyrin [48] can be regulated. In the framework of the model that we propose, this would induce changes not only in the receptor concentration on scaffold domains, but also in the size of the scaffold domain themselves, by modifying the lateral flux of scaffolding proteins onto a domain. The removal time of scaffold proteins is in our model another important parameter that could be acted upon to regulate scaffold domain size. In this context, it can be noted that smaller gephyrin domain sizes were reported in one-week old organotypic cultures as compared to four-weeks old cultures [10] while, correlatively, the mean residence time of gephyrin in scaffold domain was measured to be shorter in one-week old cultures than in four-weeks old cultures [10]. We are confident that a biophysically-rooted model such as the one we propose here offers new perspectives on synaptic dynamics, homeostasis and plasticity that it will be useful to explore in future works.

Materials and methods

Simulations

As described in the Results, we implemented two versions of the basic model, which correspond to two limits of particle rearrangement within clusters. In model A, particles within a given cluster immediately rearrange into a disc with a radius given by the number of particles n and a typical particle number density in clusters ρ, $R = \sqrt{n/(\pi\rho)}$. In model B, individual particles do not rearrange, and clusters form fractal-like structures. Both models essentially produce the same distribution of cluster sizes, see Fig 2B. Except explicitly stated otherwise, all simulation results presented were obtained with model A for reasons of reduced computational complexity.

In the simulations, all relevant quantities and parameters are expressed in the units of length $l_0 = a$ and of time $t_0 = a^2/D_0$, respectively, where a is the diameter and D_0 the diffusion constant of an individual particle. Initially, $N = c_0 L^2$ particles are randomly distributed in space in a square of side length L, and we use periodic boundary conditions throughout.

Particle diffusion and turnover are approximated in discrete time steps $\Delta t = 0.02$ small compared to the characteristic time t_0 of particle diffusion. During each step, we first update particle and cluster positions, indexed by $i = 1, \ldots, M$, by random increments according to the respective diffusion constant D_i, i.e., x and y increments drawn from a normal distribution with variance $2\Delta t D_i$ and vanishing mean. Particles or clusters that overlap afterwards fuse to give a cluster of size $n_i + n_j$, where n_i and n_j are the sizes of the fusing clusters/particles, and an accordingly updated radius. Eventually, we draw the number of desorbed particles per cluster within one time step from a binomial distribution with n_i independent draws each having probability $\Delta t\, k$ and reduce the cluster sizes and radii accordingly. For simplicity, we keep the total number of particles in the simulation constant and insert as many new, randomly distributed particles as were removed.

Simulations were generally carried out with $N \simeq 10^4$ particles, with the box size L varying correspondingly for different concentrations c_0. In our simulations, we chose a cluster density of $\rho = 0.77a^{-2}$ that corresponds to the packing fraction of a hexagonal arrangement; however, the comparison with simulations of fractal clusters shows that the results do not depend much on this choice.

Reagents

Unless otherwise noted, all reagents were purchased from Sigma-Aldrich (St. Louis, MO) or Life Technologies/Molecular Probes (Carlsbad, CA).

Lentivirus

A lentivirus encoding mEos2-Gephyrin was produced by cotransfecting the lentiviral backbone plasmid (FUGW) encoding the mEos2-Gephyrin construct (5 μg) along with the pMD2. G envelope (5 μg) and pCMVR8.74 packaging (7.5 μg) plasmids (Addgene, Cambridge, MA) into HEK293T cells using lipofectamine 2000 (60 μl). Transfection was performed in 10 cm plates once the cells reached 80% confluence. Supernatant containing lentivirus was collected 48 h after transfection, filtered through a 0.45 μm-pore-size filter, aliquoted, and stored at $-80°$C.

Cell culture and infection

All experiments were performed on dissociated spinal cord neuron cultures prepared from Sprague-Dawley rats (at E14). Experiments were carried out in accordance with the European Communities Council Directive 2010/63EU of 22 September 2010 on the protection of animals used for scientific purpose and our protocols were approved by the Charles Darwin committee in Animal experiment (Ce5/2012/018). Neurons were plated at a density of 6.3×10^4 cells/cm^2 on 18 mm coverslips pre-coated with 70 μg/ml poly-D,L-ornithine and 5% fetal calf serum. Cultures were maintained in neurobasal medium containing B-27, 2 mM glutamine, 5 U/ml penicillin, and 5 μg/ml streptomycin at 37°C and 5% CO_2. Neurons were infected at 7 days in vitro (DIV) with a recombinant lentiviral vector expressing the mEos2-Gephyrin construct.

PALM imaging

Photoactivated localization microscopy (PALM) was performed at DIV $14 - 17$ on neurons fixed with 4% paraformaldehyde and 1% sucrose (10 min). All imaging experiments were performed on an inverted Nikon Eclipse Ti microscope with a 100× oil-immersion objective (N. A. 1.49), an additional 1.5× lens, and an Andor iXon EMCCD camera. Super-resolution

movies were acquired at 20 ms frame rate under continuous illumination with activation (405 nm) and excitation (561 nm) lasers for a total of 20000 frames (6.7 minutes). Activation density was kept steady by manually increasing the activation laser intensity over time. Conventional fluorescence imaging was performed with a mercury lamp and filter sets for the detection of preconverted mEos2 (excitation 485/20, emission 525/30). The z-position was maintained during acquisition by a Nikon perfect focus system.

Single molecule detection and PALM reconstruction

The x and y coordinates of single molecule detections from each image frame were determined using an adapted version of the multiple-target tracking algorithm [49] as described previously [50]. The point spread function (PSF) signals emitted by single fluorophores were fit with a 2D Gaussian distribution. Drift in the x/y plane was corrected by calculating the relative movement of the center of mass of multiple synaptic gephyrin clusters (more than 4 per field of view), throughout the acquisition, using a sliding window of $3000 - 6000$ frames. Activations were clustered with using a single-link clustering, with a minimal distance of 50nm. Single temporal bursts in low density regions were measured for location accuracy and double counting, allowing for the counting of molecules within clusters. Regions of interests were selected by hand from fluorescent images such that only clusters within dendrites were selected for size concentration analysis. Analysis was performed on ($n = 3$) cultures.

Fits of simulated cluster size distribution to experimental data

To compare the experimentally determined cluster size distributions to the predicted distributions for a given set of parameters S, we determined the likelihood of the data according to

$$\mathcal{L}(\{\#_i\}; S) = \prod_{i=2}^{i_{\max}} \frac{n_i^{\#_i} e^{-n_i}}{\#_i!} \,, \tag{4}$$

where $n_i = A_{\exp} c_i(S)$ is the predicted count of clusters of size i in the observed surface area $A_{\exp}$ for given parameters S, and $\#_i$ is the actually determined count. The likelihood is computed including sizes up to the largest observed cluster size $i_{\max}$. For each culture ($n = 3$), we determined the likelihood over a grid of parameter values, varying turnover rate, diffusion size-dependence exponent, and total concentration. Assuming that the latter may vary between cultures, we determined the joint likelihood $\mathcal{L}_{\mathrm{tot}} = \prod_j \mathcal{L}_j$ over all three cultures $j = 1, 2, 3$ for parameters (k, σ), choosing the most likely c_0 for each (k, σ) independently for each culture. Maximum-likelihood fits of the individual clusters are shown in S3 Fig. In order to compare the experimental and simulation results, we adimensionnalized the concentrations by the density of scaffold particles in the clusters, ($\rho = 5000/3 \ \mu\mathrm{m}^{-2}$) for experiments and ($\rho = 0.77 a^{-2}$) for simulations, respectively. We assumed that gephyrin exists predominantly in trimer form and thus considered gephyrin trimers as the single-particle element in the cluster size count. We furthermore restricted the comparison of predicted and actual distributions to cluster sizes of $i \geq 2$, as counts of smaller clusters are more affected by impurities and measurement noise. The predicted particle concentrations were determined from the stationary solutions to the rate equations, see S1 Text: Rate Equations. Confidence regions corresponding to 95% probability of the optimal fit parameters were obtained by a bootstrapping technique using repeated resampling ($N = 10000$) of the observed clusters for each of the three cultures, and applying the above fit procedure to the resampled cluster size distributions.

Supporting information

S1 Text. Supporting analytical calculations. The text file includes the sections "Single scaffold domain" and "Rate equations".
(PDF)

S1 Fig. Rate-equation description of scaffold-cluster aggregation. Comparison and match of the cluster size distributions obtained from the rate-equation description to the full simulations.
(EPS)

S2 Fig. Fluorophore detections per cluster. Raw data of the fluorophore detections from which the gephyrin cluster sizes were obtained.
(EPS)

S3 Fig. Maximum-likelihood fits for individual cultures. Fits of data and confidence regions using bootstrap resampling ($N = 10000$) performed independently for each culture.
(EPS)

Acknowledgments

Author Contributions

Conceptualization: JR LGA PCR AT VH.

Investigation: JR LGA PCR AT VH.

Writing – original draft: JR LGA PCR AT VH.

References

1. Choquet D, Triller A. The Dynamic Synapse. Neuron. 2013; 80(3):691–703. https://doi.org/10.1016/j.neuron.2013.10.013 PMID: 24183020

2. Crick F. Neurobiology: Memory and molecular turnover. Nature. 1984; 312:101. https://doi.org/10.1038/312101a0

3. Rosenberg M, Meier J, Triller A, Vannier C. Dynamics of glycine receptor insertion in the neuronal plasma membrane. Journal of Neuroscience. 2001; 21(14):5036–5044. PMID: 11438579

4. Thomas P, Mortensen M, Hosie AM, Smart TG. Dynamic mobility of functional GABAA receptors at inhibitory synapses. Nature Neuroscience. 2005; 8(7):889–897. https://doi.org/10.1038/nn1483 PMID: 15951809

5. Bogdanov Y, Michels G, Armstrong-Gold C, Haydon PG, Lindstrom J, Pangalos M, et al. Synaptic GABAA receptors are directly recruited from their extrasynaptic counterparts. The EMBO Journal. 2006; 25(18):4381–4389. https://doi.org/10.1038/sj.emboj.7601309 PMID: 16946701

6. Specht CG, Izeddin I, Rodriguez PC, El Beheiry M, Rostaing P, Darzacq X, et al. Quantitative nanoscopy of inhibitory synapses: counting gephyrin molecules and receptor binding sites. Neuron. 2013; 79(2):308–321. https://doi.org/10.1016/j.neuron.2013.05.013 PMID: 23889935

7. Kim EY, Schrader N, Smolinsky B, Bedet C, Vannier C, Schwarz G, et al. Deciphering the structural framework of glycine receptor anchoring by gephyrin. The EMBO Journal. 2006; 25(6):1385–1395. https://doi.org/10.1038/sj.emboj.7601029 PMID: 16511563

8. Tyagarajan SK, Fritschy JM. Gephyrin: a master regulator of neuronal function. Nat Rev Neurosci. 2014; 15(3):141–156. https://doi.org/10.1038/nrn3670 PMID: 24552784

9. Calamai M, Specht CG, Heller J, Alcor D, Machado P, Vannier C, et al. Gephyrin Oligomerization Controls GlyR Mobility and Synaptic Clustering. Journal of Neuroscience. 2009; 29(24):7639–7648. https://doi.org/10.1523/JNEUROSCI.5711-08.2009 PMID: 19535575

10. Vlachos A, Reddy-Alla S, Papadopoulos T, Deller T, Betz H. Homeostatic regulation of gephyrin scaffolds and synaptic strength at mature hippocampal GABAergic postsynapses. Cerebral Cortex. 2013; 23(11):2700–2711. https://doi.org/10.1093/cercor/bhs260 PMID: 22918984

11. Okabe S. Fluorescence imaging of synapse formation and remodeling. Microscopy (Oxford, England). 2013; 62(1):51–62.

12. Schrader N, Kim EY, Winking J, Paulukat J, Schindelin H, Schwarz G. Biochemical characterization of the high affinity binding between the glycine receptor and gephyrin. Journal of Biological Chemistry. 2004; 279(18):18733–18741. https://doi.org/10.1074/jbc.M311245200 PMID: 14976213

13. Sola M, Bavro VN, Timmins J, Franz T, Ricard-Blum S, Schoehn G, et al. Structural basis of dynamic glycine receptor clustering by gephyrin. The EMBO Journal. 2004; 23(13):2510–2519. https://doi.org/10.1038/sj.emboj.7600256 PMID: 15201864

14. Belaidi AA, Schwarz G. Metal insertion into the molybdenum cofactor: product–substrate channelling demonstrates the functional origin of domain fusion in gephyrin. Biochemical Journal. 2013; 450 (1):149–157. https://doi.org/10.1042/BJ20121078 PMID: 23163752

15. Herweg J, Schwarz G. Splice-specific glycine receptor binding, folding, and phosphorylation of the scaffolding protein gephyrin. Journal of Biological Chemistry. 2012; 287(16):12645–12656. https://doi.org/10.1074/jbc.M112.341826 PMID: 22351777

16. Ehrensperger MV, Hanus C, Vannier C, Triller A, Dahan M. Multiple association states between glycine receptors and gephyrin identified by SPT analysis. Biophysical Journal. 2007; 92(10):3706–3718. https://doi.org/10.1529/biophysj.106.095596 PMID: 17293395

17. Ribrault C, Sekimoto K, Triller A. From the stochasticity of molecular processes to the variability of synaptic transmission. Nature Reviews Neuroscience. 2011; 12(7):375–387. https://doi.org/10.1038/nrn3025 PMID: 21685931

18. Haselwandter CA, Calamai M, Kardar M, Triller A, da Silveira RA. Formation and stability of synaptic receptor domains. Physical Review Letters. 2011; 106(23):238104. https://doi.org/10.1103/PhysRevLett.106.238104 PMID: 21770547

19. Haselwandter CA, Kardar M, Triller A, da Silveira RA. Self-assembly and plasticity of synaptic domains through a reaction-diffusion mechanism. Physical Review E. 2015; 92(3):032705. https://doi.org/10.1103/PhysRevE.92.032705

20. Burlakov V, Emptage N, Goriely A, Bressloff P. Synaptic bistability due to nucleation and evaporation of receptor clusters. Physical Review Letters. 2012; 108(2):028101. https://doi.org/10.1103/PhysRevLett.108.028101 PMID: 22324711

21. Turner MS, Sens P, Socci ND. Nonequilibrium raftlike membrane domains under continuous recycling. Physical review letters. 2005; 95(16):168301. https://doi.org/10.1103/PhysRevLett.95.168301 PMID: 16241845

22. Gómez J, Sagués F, Reigada R. Actively maintained lipid nanodomains in biomembranes. Physical Review E. 2008; 77(2):021907. https://doi.org/10.1103/PhysRevE.77.021907

23. Quang BAT, Mani M, Markova O, Lecuit T, Lenne PF. Principles of E-cadherin supramolecular organization in vivo. Current Biology. 2013; 23(22):2197–2207. https://doi.org/10.1016/j.cub.2013.09.015

24. Shouval HZ. Clusters of interacting receptors can stabilize synaptic efficacies. Proceedings of the National Academy of Sciences of the United States of America. 2005; 102(40):14440–14445. https://doi.org/10.1073/pnas.0506934102 PMID: 16189022

25. White WH. On the form of steady-state solutions to the coagulation equations. Journal Of Colloid And Interface Science. 1982; 87(1):204–208. https://doi.org/10.1016/0021-9797(82)90382-4

26. Rácz Z. Aggregation in the presence of sources and sinks: A scaling theory. Physical Review A. 1985; 32(2):1129–1133. https://doi.org/10.1103/PhysRevA.32.1129

27. Nusser Z, Roberts J, Baude A, Richards JG, Somogyi P. Relative densities of synaptic and extrasynaptic GABAA receptors on cerebellar granule cells as determined by a quantitative immunogold method. The Journal of Neuroscience. 1995; 15(4):2948–2960. PMID: 7722639

28. Saffman P, Delbrück M. Brownian motion in biological membranes. Proceedings of the National Academy of Sciences. 1975; 72(8):3111–3113. https://doi.org/10.1073/pnas.72.8.3111

29. Gambin Y, Lopez-Esparza R, Reffay M, Sierecki E, Gov N, Genest M, et al. Lateral mobility of proteins in liquid membranes revisited. Proceedings of the National Academy of Sciences of the United States of America. 2006; 103(7):2098–2102. https://doi.org/10.1073/pnas.0511026103 PMID: 16461891

30. Meakin P. Formation of fractal clusters and networks by irreversible diffusion-limited aggregation. Physical Review Letters. 1983; 51(13):1119–1122. https://doi.org/10.1103/PhysRevLett.51.1119

31. Smoluchowski Mv. Grundriß der Koagulationskinetik kolloider Lösungen. Colloid & Polymer Science. 1917; 21(3):98–104.

32. Ziff RM, McGrady E, Meakin P. On the validity of Smoluchowski's equation for cluster–cluster aggregation kinetics. The Journal of Chemical Physics. 1985; 82(11):5269–5274. https://doi.org/10.1063/1.448600

33. Dumoulin A, Levi S, Riveau B, Gasnier B, Triller A. Formation of mixed glycine and GABAergic synapses in cultured spinal cord neurons. European Journal of Neuroscience. 2000; 12(11):3883–3892. https://doi.org/10.1046/j.1460-9568.2000.00271.x PMID: 11069583

34. Pomeau Y. Symétrie des fluctuations dans le renversement du temps. Journal de Physique. 1982; 43 (6):859–867. https://doi.org/10.1051/jphys:01982004306085900

35. Pomeau Y, Le Berre M, Ginibre J. Ultimate statistical physics: fluorescence of a single atom. Journal of Statistical Mechanics: Theory and Experiment. 2016; 2016(10):104002. https://doi.org/10.1088/1742-5468/2016/10/104002

36. Hyman AA, Weber CA, Jülicher F. Liquid-liquid phase separation in biology. Annual Review of Cell and Developmental Biology. 2014; 30:39–58. https://doi.org/10.1146/annurev-cellbio-100913-013325 PMID: 25288112

37. Zeng M, Shang Y, Araki Y, Guo T, Huganir RL, Zhang M. Phase Transition in Postsynaptic Densities Underlies Formation of Synaptic Complexes and Synaptic Plasticity. Cell. 2016; 166(5):1163–1175. https://doi.org/10.1016/j.cell.2016.07.008 PMID: 27565345

38. Brangwynne CP. Phase transitions and size scaling of membrane-less organelles. The Journal of Cell Biology. 2013; 203(6):875–881. https://doi.org/10.1083/jcb.201308087 PMID: 24368804

39. Singer JH, Berger AJ. Contribution of single-channel properties to the time course and amplitude variance of quantal glycine currents recorded in rat motoneurons. Journal of Neurophysiology. 1999; 81 (4):1608–1616. PMID: 10200197

40. Rigo JM, Badiu CI, Legendre P. Heterogeneity of postsynaptic receptor occupancy fluctuations among glycinergic inhibitory synapses in the zebrafish hindbrain. The Journal of Physiology. 2003; 553(Pt 3):819–832. https://doi.org/10.1113/jphysiol.2003.049577 PMID: 14500774

41. Witten TA, Sander LM. Diffusion-limited aggregation, a kinetic critical phenomenon. Physical Review Letters. 1981; 47(19):1400–1403. https://doi.org/10.1103/PhysRevLett.47.1400

42. Kessler DA, Koplik J, Levine H. Pattern selection in fingered growth phenomena. Advances in Physics. 1988; 37(3):255–339. https://doi.org/10.1080/00018738800101379

43. Minerbi A, Kahana R, Goldfeld L, Kaufman M, Marom S, Ziv NE. Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity. PLoS Biol. 2009; 7(6):e1000136. https://doi.org/10.1371/journal.pbio.1000136 PMID: 19554080

44. Kesten H. Random difference equations and renewal theory for products of random matrices. Acta Mathematica. 1973; 131(1):207–248. https://doi.org/10.1007/BF02392040

45. Statman A, Kaufman M, Minerbi A, Ziv NE, Brenner N. Synaptic size dynamics as an effectively stochastic process. PLoS Comput Biol. 2014; 10(10):e1003846. https://doi.org/10.1371/journal.pcbi.1003846 PMID: 25275505

46. Rubinski A, Ziv NE. Remodeling and Tenacity of Inhibitory Synapses: Relationships with Network Activity and Neighboring Excitatory Synapses. PLoS Comput Biol. 2015; 11(11):e1004632. https://doi.org/10.1371/journal.pcbi.1004632 PMID: 26599330

47. Tatavarty V, Sun Q, Turrigiano GG. How to scale down postsynaptic strength. The Journal of Neuroscience. 2013; 33(32):13179–13189. https://doi.org/10.1523/JNEUROSCI.1676-13.2013 PMID: 23926271

48. Zita MM, Marchionni I, Bottos E, Righi M, Del Sal G, Cherubini E, et al. Post-phosphorylation prolyl isomerisation of gephyrin represents a mechanism to modulate glycine receptors function. The EMBO Journal. 2007; 26(7):1761–1771. https://doi.org/10.1038/sj.emboj.7601625 PMID: 17347650

49. Sergé A, Bertaux N, Rigneault H, Marguet D. Dynamic multiple-target tracing to probe spatiotemporal cartography of cell membranes. Nature Methods. 2008; 5(8):687–694. https://doi.org/10.1038/nmeth.1233 PMID: 18604216

50. Izeddin I, Specht CG, Lelek M, Darzacq X, Triller A, Zimmer C, et al. Super-resolution dynamic imaging of dendritic spines using a low-affinity photoconvertible actin probe. PloS one. 2011; 6(1):e15611. https://doi.org/10.1371/journal.pone.0015611 PMID: 21264214

Cerebellar learning using perturbations

Guy Bouvier[1†‡], Johnatan Aljadeff[2§], Claudia Clopath[3], Célian Bimbard[1#], Jonas Ranft[1], Antonin Blot[1¶], Jean-Pierre Nadal[4,5], Nicolas Brunel[2], Vincent Hakim[4], Boris Barbour[1*]**

[1]Institut de biologie de l'École normale supérieure (IBENS), École normale supérieure, CNRS, INSERM, PSL University, Paris, France; [2]Departments of Statistics and Neurobiology, University of Chicago, Chicago, United States; [3]Department of Bioengineering, Imperial College London, London, United Kingdom; [4]Laboratoire de Physique Statistique, École normale supérieure, CNRS, PSL University, Sorbonne Université, Paris, France; [5]Centre d'Analyse et de Mathématique Sociales, EHESS, CNRS, PSL University, Paris, France

*For correspondence:
boris.barbour@ens.fr

Present address: [†]Department of Physiology, University of California, San Francisco, San Francisco, United States; [‡]Sandler Neuroscience, University of California, San Francisco, San Francisco, United States; [§]Department of Bioengineering, Imperial College London, London, United Kingdom; [#]Laboratoire des systèmes perceptifs, Département d'études cognitives, École normale supérieure, PSL University, Paris, France; [¶]Sainsbury-Wellcome Centre for Neural Circuits and Behaviour, University College London, London, United Kingdom; [**]Departments of Neurobiology and Physics, Duke University, Durham, United States

Abstract The cerebellum aids the learning of fast, coordinated movements. According to current consensus, erroneously active parallel fibre synapses are depressed by complex spikes signalling movement errors. However, this theory cannot solve the *credit assignment problem* of processing a global movement evaluation into multiple cell-specific error signals. We identify a possible implementation of an algorithm solving this problem, whereby spontaneous complex spikes perturb ongoing movements, create eligibility traces and signal error changes guiding plasticity. Error changes are extracted by adaptively cancelling the average error. This framework, *stochastic gradient descent with estimated global errors* (SGDEGE), predicts synaptic plasticity rules that apparently contradict the current consensus but were supported by plasticity experiments in slices from mice under conditions designed to be physiological, highlighting the sensitivity of plasticity studies to experimental conditions. We analyse the algorithm's convergence and capacity. Finally, we suggest SGDEGE may also operate in the basal ganglia.
DOI: https://doi.org/10.7554/eLife.31599.001

Introduction

A central contribution of the cerebellum to motor control is thought to be the learning and automatic execution of fast, coordinated movements. Anatomically, the cerebellum consists of a convoluted, lobular cortex surrounding the cerebellar nuclei (*Figure 1A*, *Eccles et al., 1967*; *Ito, 1984*). The main input to the cerebellum is the heterogeneous mossy fibres, which convey multiple modalities of sensory, contextual and motor information. They excite both the cerebellar nuclei and the cerebellar cortex; in the cortex they synapse with the very abundant granule cells, whose axons, the parallel fibres, excite Purkinje cells. Purkinje cells constitute the sole output of the cerebellar cortex and project an inhibitory connection to the nuclei, which therefore combine a direct and a transformed mossy fibre input with opposite signs. The largest cell type in the nuclei, the projection neurone, sends excitatory axons to several motor effector systems, notably the motor cortex via the thalamus. Another nuclear cell type, the nucleo-olivary neurone, inhibits the inferior olive. The cerebellum receives a second external input: climbing fibres from the inferior olive, which form an extensive, ramified connection with the proximal dendrites of the Purkinje cell. Each Purkinje cell receives a single climbing fibre. A more modular diagram of the olivo-cerebellar connectivity relevant to this paper is shown in *Figure 1B*; numerous cell types and connections have been omitted for simplicity.

Purkinje cells discharge two distinct types of action potential (*Figure 1C*). They nearly continuously emit *simple spikes*—standard, if brief, action potentials—at frequencies that average 50 Hz. This frequency is modulated both positively and negatively by the intensity of inputs from the mossy
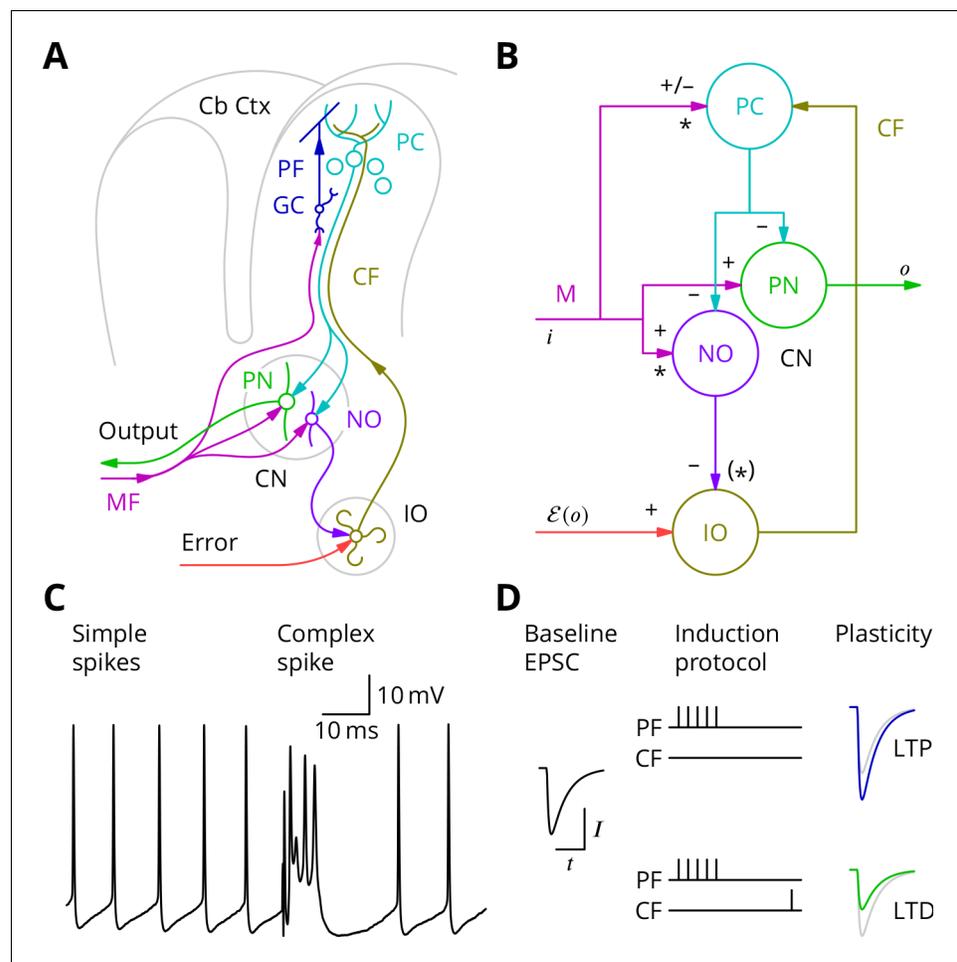
Figure 1. The cerebellar circuitry and properties of Purkinje cells. (**A**) Simplified circuit diagram. MF, mossy fibres; CN, (deep) cerebellar nuclei; GC, granule cells; Cb Ctx, cerebellar cortex; PF, parallel fibres; PC, Purkinje cells; PN, projection neurones; NO, nucleo-olivary neurones; IO, inferior olive; CF, Climbing fibres. (**B**) Modular diagram. The ± signs next to the synapses indicate whether they are excitatory or inhibitory. The granule cell and indirect inhibitory inputs they recruit have been subsumed into a bidirectional mossy fibre–Purkinje cell input, M. Potentially plastic inputs of interest here are denoted with an asterisk. i, input; o, output; $\mathcal{E}(o)$, error (which is a function of the output). (**C**) Typical Purkinje cell electrical activity from an intracellular patch-clamp recording. Purkinje cells fire two types of action potential: simple spikes and, in response to climbing fibre input, complex spikes. (**D**) According to the consensus plasticity rule, a complex spike will depress parallel fibre synapses active about 100 ms earlier. The diagram depicts idealised excitatory postsynaptic currents (EPSCs) before and after typical induction protocols inducing long-term potentiation (LTP) or depression (LTD). Grey, control EPSC; blue, green, post-induction EPSCs.
DOI: https://doi.org/10.7554/eLife.31599.002

fibre–granule cell pathway (which can also recruit interneurons that inhibit Purkinje cells; *Eccles et al., 1967*). Such modulations of Purkinje cell firing are thought to underlie their contributions to motor control. In addition, when the climbing fibre is active, an event that occurs continuously but in a somewhat irregular pattern with a mean frequency of around 1 Hz, the Purkinje cell emits a completely characteristic *complex spike* under the influence of the intense excitation from the climbing fibre (*Figure 1C*).

The history of research into cerebellar learning is dominated by the theory due to *Marr (1969)* and *Albus (1971)*. They suggested that the climbing fibre acts as a 'teacher' to guide plasticity of parallel fibre–Purkinje cell synapses. It was several years, however, before experimental support for this hypothesis was obtained (*Ito et al., 1982*; *Ito and Kano, 1982*), by which time the notion that

the climbing fibre signalled errors had emerged (*Ito, 1972*; *Ito, 1984*). Error modalities thought to be represented by climbing fibres include: pain, unexpected touch, imbalance, and retinal slip. According to the modern understanding of this theory, by signalling such movement errors, climbing fibres induce long-term depression (LTD) of parallel fibre synapses that were active at the same time (*Ito et al., 1982*; *Ito and Kano, 1982*; *Sakurai, 1987*; *Crepel and Jaillard, 1991*) or, more precisely, shortly before (*Wang et al., 2000*; *Sarkisov and Wang, 2008*; *Safo and Regehr, 2008*). A compensating long-term potentiation (LTP) is necessary to prevent synaptic saturation (*Lev-Ram et al., 2002*; *Lev-Ram et al., 2003*; *Coesmans et al., 2004*) and its induction is reported to follow high-frequency parallel fibre activity in the absence of complex spikes (*Jörntell and Ekerot, 2002*; *Bouvier et al., 2016*). Plasticity of parallel fibre synaptic currents according to these plasticity rules is diagrammed in *Figure 1D*.

Cerebellar learning with the Marr-Albus-Ito theory has mostly been considered, both experimentally and theoretically, at the level of single cells or of uniformly responding groups of cells learning a single stereotyped adjustment. Such predictable and constrained movements, exemplified by eye movements and simple reflexes, provide some of the best studied models of cerebellar learning: the vestibulo-ocular reflex (*Robinson, 1976*; *Ito et al., 1974*; *Blazquez et al., 2004*), nictitating membrane response/eye blink conditioning (*McCormick et al., 1982*; *Yeo et al., 1984*; *Yeo and Hesslow, 1998*), saccade adaptation (*Optican and Robinson, 1980*; *Dash and Thier, 2014*; *Soetedjo et al., 2008*) and regulation of limb movements by withdrawal reflexes (*Ekerot et al., 1995*; *Garwicz et al., 2002*). All of these motor behaviours have in common that there could conceivably be a fixed mapping between an error and a suitable corrective action. Thus, adaptations necessary to ensure gaze fixation are exactly determined by the retinal slip. Such fixed error-correction relations may have been exploited during evolution to create optimised correction circuitry.

The problems arise with the Marr-Albus-Ito theory if one tries to extend it to more complex situations, where neurones must respond heterogeneously (not uniformly) and/or where the flexibility to learn arbitrary responses is required. Many motor control tasks, for instance coordinated movements involving the hands, can be expected to fall into this class. To learn such complex/arbitrary movements with the Marr-Albus-Ito theory requires error signals that are specific for each cell, each movement and each time within each movement. The theory is thus incomplete, because it does not describe how a global evaluation of movement error can be processed to provide such detailed instructions for plasticity to large numbers of cells, a general difficulty which the brain has to face that was termed the *credit assignment problem* by *Minsky (1961)*.

A suitable algorithm for solving the general cerebellar learning problem would be *stochastic gradient descent*, a classical optimisation method that *Minsky (1961)* suggested might operate in the brain. We shall speak of 'descent' when minimising an error and 'ascent' when maximising a reward, but the processes are mathematically equivalent. In stochastic gradient descent, the objective function is explored by random variations in the network that alter behaviour, with plasticity then retaining those variations that improve the behaviour, as signalled by a decreased error or increased reward. Several possible mechanisms of varying biological plausibility have been proposed. In particular, perturbations caused by synaptic release (*Minsky, 1954*; *Seung, 2003*) or external inputs (*Doya and Sejnowski, 1988*) have been suggested, while various (abstract) mechanisms have been proposed for extraction of changes in the objective function (*Williams, 1992*). To avoid confusion, note that these forms of stochastic gradient descent differ from those conforming to the more restrictive definition used in the machine learning community, in which the stochastic element is the random sampling of examples from the training set (*Robbins and Monro, 1951*; *Shalev-Shwartz and Ben-David, 2014*). This latter form has achieved broad popularity in online learning from large data sets and in deep learning. Although the theoretical framework for (perturbative) stochastic gradient descent is well established, the goal of identifying in the brain a network and cellular implementation of such an algorithm has proved elusive.

The learning behaviour with the best established resemblance to stochastic gradient ascent is the acquisition of song in male songbirds. The juvenile song is refined by a trial and error process to approach a template memorised from a tutor during a critical period (*Konishi, 1965*; *Mooney, 2009*). The analogy with stochastic gradient ascent was made by *Doya and Sejnowski (1988)* and was then further developed experimentally (*Olveczky et al., 2005*) and theoretically (*Fiete et al., 2007*). However, despite these very suggestive behavioural correlates, relatively little

progress has been made in verifying model predictions for plasticity or identifying the structures responsible for storing the template and evaluating its match with the song.

A gradient descent mechanism for the cerebellum has been proposed by the group of *Dean et al. (2002)*, who term their algorithm *decorrelation*. Simply stated (*Dean and Porrill, 2014*), if both parallel fibre and climbing fibre inputs to a Purkinje cell are assumed to vary about their respective mean values, their correlation or anticorrelation indicates the local gradient of the error function and thus the sign of the required plasticity. At the optimum, which is a minimum of the error function, there should be no correlation between variations of the climbing fibre rate and those of the parallel fibre input. Hence the name: the algorithm aims to decorrelate parallel and climbing fibre variations. An appropriate plasticity rule for implementing decorrelation is a modified covariance rule (*Sejnowski, 1977*; who moreover suggested in abridged form a similar application to cerebellar learning). Although decorrelation provides a suitable framework, its proponents are still in the process of developing a cellular implementation (*Menzies et al., 2010*). Moreover, we believe that the detailed implementation suggested by *Menzies et al., 2010* is unable to solve a temporal credit assignment problem (that we identify below) arising in movements that can only be evaluated upon completion.

Below, we analyse in more detail how the Marr-Albus-Ito theory fails to solve the credit assignment problem and suggest how the cerebellum might implement stochastic gradient descent. We shall provide support for unexpected predictions the proposed implementation makes regarding cerebellar synaptic plasticity, which are different from those suggested for decorrelation. Finally, we shall perform a theoretical analysis demonstrating that learning with the algorithm converges and that it is able to attain maximal storage capacity.

Results

Requirements for cerebellar learning

We begin by examining the current consensus (Marr-Albus-Ito) theory of cerebellar learning and illustrating some of its limitations when extended to the optimisation of complex, arbitrary movements. The learning framework we consider is the following. The cerebellar circuitry must produce at its output trains of action potentials of varying frequencies at given times (*Thach, 1968*). We consider only firing rates $r(t)$, which are a function of time in the movement (*Figure 2A*). The cerebellar output influences the movement and the resulting movement error, which can be sensed and fed back to the Purkinje cells in the form of climbing fibre activity.

We constructed a simple network simulation to embody this framework (see Materials and methods). In it, mossy fibre drive to the cerebellum was considered to be movement- and time-dependent but not to vary during learning. Granule cells and molecular layer interneurons were collapsed into the mossy fibre inputs, which acted on Purkinje cells through unconstrained synapses (negative weights could arise through plasticity of interneuron synapses; *Jörntell and Ekerot, 2002*; *Jörntell and Ekerot, 2003*; *Mittmann and Häusser, 2007*) that were modified to optimise the firing profiles of the projection neurones of the deep cerebellar nuclei. We implemented the Marr-Albus-Ito algorithm via a plasticity rule in which synaptic weights were updated by the product of presynaptic activity and the presence of the error signal. Thus, active mossy fibre (pathway) inputs would undergo LTP (increasing Purkinje cell firing; *Lev-Ram et al., 2003*) after each movement unless the climbing fibre was active, in which case they would undergo LTD (decreasing activity). The model represented a microzone whose climbing fibres were activated uniformly by a global movement error.

The behaviour of the Marr-Albus-Ito algorithm in this network depends critically on how the definition of the error is extended from the error for an individual cell. The most natural definition for minimising all errors would be to sum the absolute differences between the actual firing profile of each cerebellar nuclear projection neurone and its target, see *Equation 7*. However, this error definition is incompatible with the Marr-Albus-Ito algorithm (*Unsigned* in *Figure 2B*). Examination in *Figure 2C* of the firing profile of one of the projection neurones shows that their rate simply saturates. The algorithm *is* able to minimise the average signed error defined in *Equation 14* (*Signed* in *Figure 2B*). However, inspection of a specimen final firing profile illustrates the obvious limitation
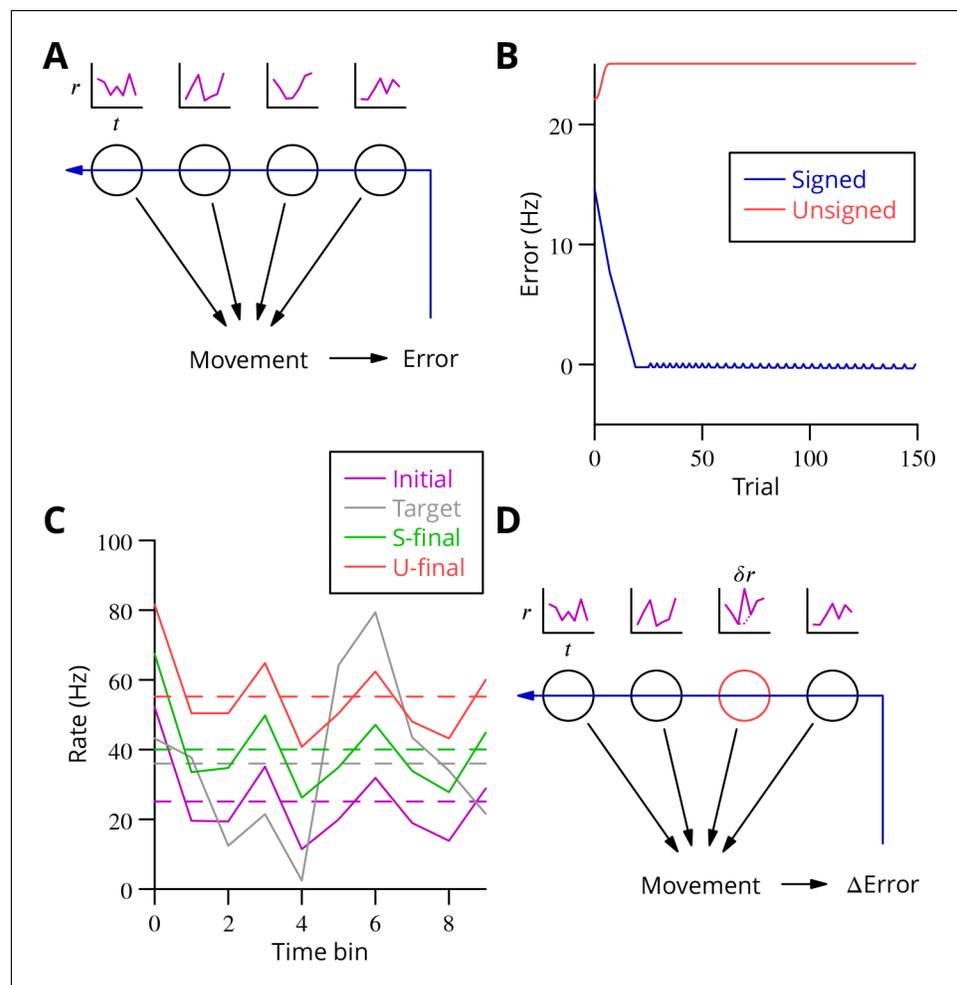
Figure 2. Analysis of cerebellar learning. (**A**) The model of cerebellar learning we address is to adjust the temporal firing profiles (miniature graphs of rate as a function of time, $r(t)$, *magenta*) of multiple cerebellar output neurones (nuclear projection neurones) (*black circles*) to optimise a movement, exploiting the evaluated movement error, which is fed back to the cerebellar circuit (Purkinje cells) (*blue arrow*). (**B**) The Marr-Albus-Ito theory was simulated in a network simulation embodying different error definitions (Results, Materials and methods). The Marr-Albus-Ito algorithm is able to minimise the average signed error (*blue, Signed*) of a group of projection neurones, but not the unsigned error (*red, Unsigned*). (**C**) However, as expected, optimising the signed error does not optimise individual cell firing profiles: comparison of the underlying firing profiles (initial, *magenta*; final, *green*) with their target (*grey*) for a specimen projection neurone illustrates that neither the temporal profile nor even the mean single-cell firing rate is optimised. If the unsigned error is used, there is no convergence and the firing rates simply saturate (*red*) and the error *increases*. (**D**) A different learning algorithm will be explored below: stochastic gradient descent, according to which temporally and spatially localised firing perturbations (δr, *red*, of a Purkinje cell) are consolidated if the movement improves; this requires extraction of the *change* of error (Δ Error).
DOI: https://doi.org/10.7554/eLife.31599.003

that neither the detail of the temporal profile nor even the mean firing rate for individual cells is optimised when using this error (see *Figure 2C*).

This limitation illustrates the credit assignment problem. How to work back from a single evaluation after the movement to optimise the firing rate of multiple cells at different time points, with all of those firing profiles also differing between multiple movements which must all be learnt in parallel? A model problem that the Marr-Albus-Ito algorithm cannot solve would be two neurones receiving the same inputs and error signal but needing to undergo opposite plasticity. In the general case of learning complex, arbitrary movements, this requires impractical foreknowledge to be embodied by the climbing fibre system, to know exactly which cell must be depressed to optimise each

movement, and this still would not solve the problem of inducing plasticity differentially at different time points during a firing profile, given that often only a single, post-movement evaluation will be available. An example would be a saccade, where the proximity to the target can only be assessed once the movement of the eye has ceased.

The complex spike as trial and error

The apparent difficulty of solving the credit assignment problem within the Marr-Albus-Ito algorithm led us to consider whether the cerebellum might implement a form of stochastic gradient descent. By combining a known perturbation of one or a small group of neurones with information about the change of the global error (*Figure 2D*), it becomes possible to consolidate, specifically in perturbed neurones, those motor command modifications that reduce the error, thus leading to a progressive optimisation. We set out to identify a biologically plausible implementation in the cerebellum. In particular, that implementation should be able to optimise multiple movements in parallel, employing only feasible cellular computations—excitation, inhibition and thresholding.

A cerebellar implementation of stochastic gradient descent must include a source of perturbations of the Purkinje cell firing rate δr. The fact that Purkinje cells can contribute to different movements with arbitrary and unknown sequencing imposes an implementation constraint preventing simple-minded approaches like comparing movements performed twice in succession. We recall that we assume that no explicit information categorising or identifying movements is available to the Purkinje cell. It is therefore necessary that knowledge of both the presence and sign of δr be available within the context of a single movement execution.

In practice, a number of different perturbation mechanisms can still satisfy these requirements. For instance, any binary signal would be suitable, since the sign of the perturbation with respect to its mean would be determined by the simple presence or absence of the signal. Several plausible mechanisms along these lines have been proposed, including external modulatory inputs (*Doya and Sejnowski, 1988*; *Fiete et al., 2007*), failures and successes of synaptic transmission (*Seung, 2003*) or the absence and presence of action potentials (*Xie and Seung, 2004*). However, none of these mechanisms has yet attracted experimental support at the cellular level.

In the cerebellar context, parallel fibre synaptic inputs are so numerous that the correlation between individual input variations and motor errors is likely to be extremely weak, whereas we seek a perturbation that is sufficiently salient to influence ongoing movement. Purkinje cell action potentials are also a poor candidate, because they are not back-propagated to parallel fibre synapses (*Stuart and Häusser, 1994*) and therefore probably cannot guide their plasticity, but the ability to establish a synaptic eligibility trace is required. Bistable firing behaviour of Purkinje cells (*Loewenstein et al., 2005*; *Yartsev et al., 2009*), with the down-state (or long pauses) representing a clear perturbation towards lower (zero) firing rates, is a perturbation candidate. However, exploratory plasticity experiments did not support this hypothesis and the existence of bistability in vivo is disputed (*Schonewille et al., 2006a*).

We thus propose, in accordance with a suggestion due to *Harris (1998)*, another possible perturbation of Purkinje cell firing: the complex spike triggered by the climbing fibre. We note that there are probably two types of inferior olivary activity. Olivary neurones mediate classical error signalling triggered by external synaptic input, but they also exhibit continuous and irregular spontaneous activity in the absence of overt errors. We suggest the spontaneous climbing fibre activations cause synchronised *perturbation complex spikes* (pCSs) in small groups of Purkinje cells via the ~ 1:10 inferior olivary–Purkinje cell divergence (*Schild, 1970*; *Mlonyeni, 1973*; *Caddy and Biscoe, 1976*), dynamic synchronisation of olivary neurones through electrical coupling (*Llinás and Yarom, 1986*; *Bazzigaluppi et al., 2012a*) and common synaptic drive. The excitatory perturbation—a brief increase of firing rate (*Ito and Simpson, 1971*; *Campbell and Hesslow, 1986*; *Khaliq and Raman, 2005*; *Monsivais et al., 2005*)—feeds through the cerebellar nuclei (changing sign; *Bengtsson et al., 2011*) to the ongoing motor command and causes a perturbation of the movement, which in turn may modify the error of the movement.

The perturbations are proposed to guide learning in the following manner. If a perturbation complex spike results in an increase of the error, the raised activity of the perturbed Purkinje cells was a mistake and reduced activity would be preferable; parallel fibre synapses active at the time of the perturbing complex spikes should therefore be depressed. Conversely, if the perturbation leads to a

reduction of error (or does not increase it), the increased firing rate should be consolidated by potentiation of the simultaneously active parallel fibres.

How could an increase of the error following a perturbation be signalled to the Purkinje cell? We suggest that the climbing fibre also performs this function. Specifically, if the perturbation complex spike increases the movement error, a secondary *error complex spike* (eCS) is emitted shortly afterwards, on a time scale of the order of 100 ms (50–300 ms). This time scale is assumed because it corresponds to the classical error signalling function of the climbing fibre, because it allows sufficient time for feedback via the error modalities known to elicit complex spikes (touch, pain, balance, vision) and because such intervals are known to be effective in plasticity protocols (*Wang et al., 2000*; *Sarkisov and Wang, 2008*; *Safo and Regehr, 2008*). The interval could also be influenced by the oscillatory properties of olivary neurones (*Llinás and Yarom, 1986*; *Bazzigaluppi et al., 2012b*).

The predicted plasticity rule is therefore as diagrammed in *Figure 3*. Only granule cell synapses active simultaneously with the perturbation complex spike undergo plasticity, with the sign of the plasticity being determined by the presence or absence of a subsequent error complex spike. Granule cell synapses active in the absence of a synchronous perturbation complex spike should not undergo plasticity, even if succeeded by an error complex spike. We refer to these different protocols with the abbreviations (and give our predicted outcome in parenthesis): G_ _ (no change), GP_ (LTP), G_E (no change), GPE (LTD), where G indicates granule cell activity, P the presence of a perturbation complex spike and E the presence of an error complex spike. Note that both granule cells and climbing fibres are likely to be active in high-frequency bursts rather than the single activations idealised in *Figure 3*.

Several of the predictions of this rule appear to be incompatible with the current consensus. Thus, parallel fibre synapses whose activity is simultaneous with (GP_) or followed by a complex spike (G_E) have been reported to be depressed (*Sakurai, 1987*; *Crepel and Jaillard, 1991*; *Lev-Ram et al., 1995*; *Coesmans et al., 2004*; *Safo and Regehr, 2008*; *Gutierrez-Castellanos et al.,*
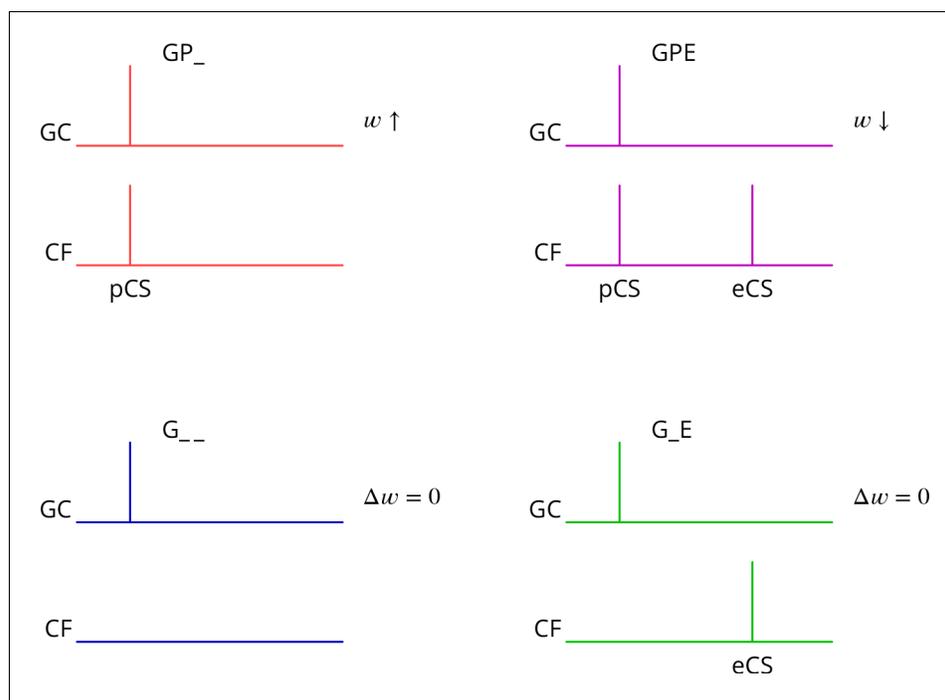
Figure 3. Predicted plasticity rules. Synchronous activation of granule cell synapses and a perturbation complex spike (pCS) leads to LTP (GP_, increased synaptic weight *w*; top left, *red*), while the addition of a succeeding error complex spike (eCS) leads to LTD (GPE, top right, *magenta*). The bottom row illustrates the corresponding 'control' cases from which the perturbation complex spike is absent; no plasticity should result (G_ _ *blue* and G_E *green*).
DOI: https://doi.org/10.7554/eLife.31599.004

2017), while we predict potentiation and no change, respectively. Furthermore, parallel fibre activity alone (G_ _) leads to potentiation (*Lev-Ram et al., 2002*; *Jörntell and Ekerot, 2002*; *Lev-Ram et al., 2003*; *Coesmans et al., 2004*; *Gutierrez-Castellanos et al., 2017*), while we predict no change.

Synaptic plasticity under physiological conditions

As described above, the plasticity rules we predict for parallel fibre–Purkinje cell synapses are, superficially at least, close to the opposite of the consensus in the literature. Current understanding of the conditions for inducing plasticity gives a key role to the intracellular calcium concentration (combined with nitric oxide signalling; *Coesmans et al., 2004*; *Bouvier et al., 2016*), whereby high intracellular calcium concentrations are required for LTD and moderate concentrations lead to LTP. Standard experimental conditions for studying plasticity in vitro, notably the extracellular concentration of calcium, are likely to result in more elevated intracellular calcium concentrations during induction than pertain physiologically. Recognising that this could alter plasticity outcomes, we set out to test whether our predicted plasticity rules might be verified under more physiological conditions.

We made several changes to standard protocols (see Materials and methods); one was cerebellum-specific, but the others also apply to in vitro plasticity studies in other brain regions. We did not block GABAergic inhibition. We lowered the extracellular calcium concentration from the standard 2 mM (or higher) used in slice work to 1.5 mM (*Pugh and Raman, 2008*), which is near the maximum values measured in vivo in rodents (*Nicholson et al., 1978*; *Jones and Keep, 1988*; *Silver and Erecińska, 1990*). To avoid the compact bundles of active parallel fibres produced by the usual stimulation in the molecular layer, we instead used weak granule cell layer stimuli, which results in sparse and spatially dispersed parallel fibre activity. Interestingly, it has been reported that standard protocols using granule cell stimulation are unable to induce LTD (*Marcaggi and Attwell, 2007*). We used a pipette solution designed to prolong energy supply in extended cells like the Purkinje cell (see Materials and methods). Experiments were carried out in adult mouse sagittal cerebellar slices using otherwise standard patch-clamp techniques.

Pairs of granule cell test stimuli with an interval of 50 ms were applied at 0.1 Hz before and after induction; EPSCs were recorded in voltage clamp at −70 mV. Pairs of climbing fibre stimuli with a 2.5 ms interval were applied at 0.5 Hz throughout the test periods, mimicking tonic climbing fibre activity, albeit at a slightly lower rate. The interleaved test granule cell stimulations were sequenced 0.5 s before the climbing fibre stimulations. The analysis inclusion criteria and amplitude measurement for the EPSCs are detailed in the Materials and methods. The average amplitude of the granule cell EPSCs retained for analysis was −62 ± 46 pA (mean ± s.d., n = 58). The rise and decay time constants (of the global averages) were 0.74 ± 0.36 ms and 7.2 ± 2.7 ms (mean ± sd), respectively.

During induction, performed in current clamp without any injected current, the granule cell input consisted of a burst of five stimuli at 200 Hz, reproducing the propensity of granule cells to burst at high frequencies (*Chadderton et al., 2004*; *Jörntell and Ekerot, 2006*). The climbing fibre input reflected the fact that these can occur in very high-frequency bursts (*Eccles et al., 1966*; *Maruta et al., 2007*). We used two stimuli at 400 Hz to represent the perturbation complex spike and four for the subsequent error complex spike if it was included in the protocol. Depending on the protocol, the climbing fibre stimuli had different timings relative to the granule cell stimuli: a pair of climbing fibre stimuli at 400 Hz, 11–15 ms or ∼500 ms after the start of the granule cell burst and/or four climbing fibre stimuli at 400 Hz, 100–115 ms after the beginning of the granule cell burst (timing diagrams will be shown in the Results). In a fraction of cells, the climbing fibre stimuli after the first were not reliable; our grounds for including these cells are detailed in the Materials and methods. The interval between the two bursts of climbing fibre stimuli when the error complex spike was present was about 100 ms. We increased the interval between induction episodes from the standard one second to two, to reduce any accumulating signal during induction. 300 such induction sequences were applied (*Lev-Ram et al., 2002*).

We first show the protocols relating to LTP (*Figure 4*). A granule cell burst was followed by a distant perturbation climbing fibre stimulus or the two inputs were activated simultaneously. In the examples shown, the protocol with simultaneous activation (GP_, *Figure 4C,D*) caused a potentiation of about 40%, while the temporally separate inputs caused a smaller change of 15% in the opposite direction (G_ _, *Figure 4A,B*). We note that individual outcomes were quite variable; group data and statistics will be shown below. The mean paired-pulse ratio in our recordings was A2/
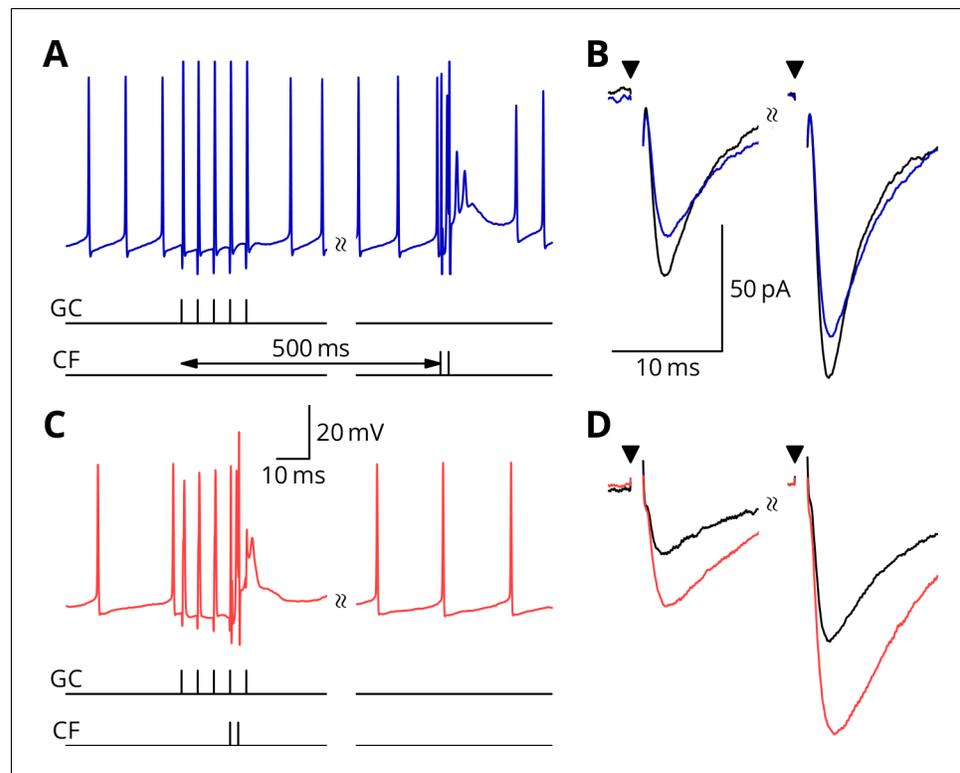
Figure 4. Simultaneous granule cell and climbing fibre activity induces LTP. (A) Membrane potential (*blue*) of a Purkinje cell during an induction protocol (G_ _) where a burst of 5 granule cell stimuli at 200 Hz was followed after 0.5 s by a pair of climbing fibre stimuli at 400 Hz. (B) Average EPSCs recorded up to 10 min before (*black*) and 20–30 min after the end of the protocol of A (*blue*). Paired test stimuli (*triangles*) were separated by 50 ms and revealed the facilitation typical of the granule cell input to Purkinje cells. In this case, the induction protocol resulted in a small reduction (*blue* vs. *black*) of the amplitude of responses to both pulses. (C) Purkinje cell membrane potential (*red*) during a protocol (GP_) where the granule cells and climbing fibres were activated simultaneously, with timing otherwise identical to A. (D) EPSCs recorded before (*black*) and after (*red*) the protocol in C. A clear potentiation was observed in both of the paired-pulse responses.
DOI: https://doi.org/10.7554/eLife.31599.005

A1 = 1.75 ± 0.32 (mean $\pm$ sd, $n = 58$). As here, no significant differences of paired-pulse ratio were observed with any of the plasticity protocols: plasticity $-$ baseline difference for GP_, mean -0.08, 95% confidence interval (-0.34, 0.20), $n = 15$; GPE, mean 0.12, 95 % c.i. (-0.07, 0.33), $n = 10$; G_ _, mean -0.01, 95 % c.i. (-0.24, 0.24), $n = 18$; G_E, mean -0.09, 95 % c.i. (-0.23, 0.29), $n = 15$.

Figure 5 illustrates tests of our predictions regarding the induction of LTD. As before, a granule cell burst was paired with the perturbation climbing fibre, but now a longer burst of climbing fibre stimuli was appended 100 ms later, representing an error complex spike (GPE, *Figure 5C,D*). A clear LTD of about 40% developed following the induction. In contrast, if the perturbation complex spike was omitted, leaving the error complex spike (G_E, *Figure 5A,B*), no clear change of synaptic weight occurred (an increase of about 10%). During induction, cells would generally begin in a tonic firing mode, but nearly all ceased firing by the end of the protocol. The specimen sweeps in *Figure 5* are taken towards the end of the induction period and illustrate the Purkinje cell responses when spiking had ceased.

The time courses of the changes of EPSC amplitude are shown in normalised form in *Figure 6* (see Materials and methods). The individual data points of the relative EPSC amplitudes for the different protocols are also shown. A numerical summary of the group data and statistical comparisons is given in *Table 1*.

In a complementary series of experiments, we explored the plasticity outcome when six climbing fibre stimuli were grouped in a burst and applied simultaneously with the granule cell burst. This
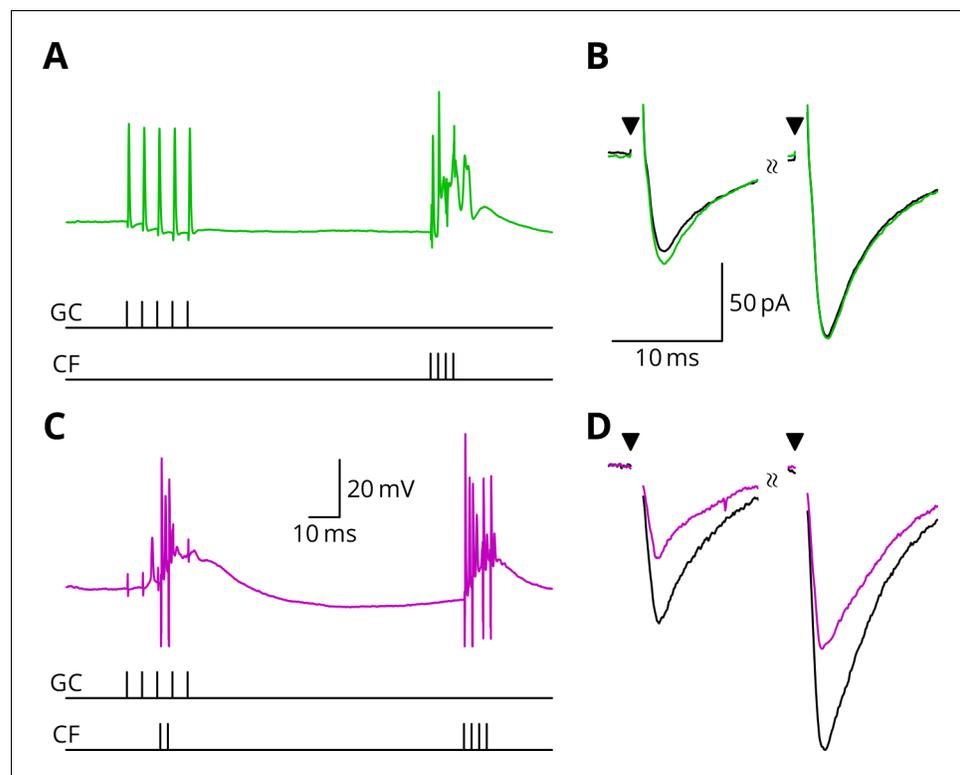
Figure 5. LTD requires simultaneous granule cell and climbing fibre activity closely followed by an additional complex spike. (**A**) Membrane potential of a Purkinje cell (*green*) during a protocol where a burst of five granule cell stimuli at 200 Hz was followed after 100 ms by four climbing fibre stimuli at 400 Hz (G_E). (**B**) Average EPSCs recorded up to 10 min before (*black*) and 20–30 min after the end of the protocol of A (*green*). The interval between the paired test stimuli (*triangles*) was 50 ms. The induction protocol resulted in little change (*green* vs. *black*) of the amplitude of either pulse. (**C**) Purkinje cell membrane potential (*magenta*) during the same protocol as in A with the addition of a pair of climbing fibre stimuli simultaneous with the granule cell stimuli (GPE). (**D**) EPSCs recorded before (*black*) and after (*magenta*) the protocol in C. A clear depression was observed.
DOI: https://doi.org/10.7554/eLife.31599.006

allowed comparison with the GPE protocol, which also contained 4 + 2 = 6 climbing fibre stimuli. The results in *Figure 6—figure supplement 1* shows that a modest LTP was observed on average: after/before ratio 1.12 $\pm$ 0.12 (mean $\pm$ SEM); 95% c.i. 0.89–1.35. This result is clearly different from the LTD observed under GPE (p=0.0034; two-tailed Wilcoxon rank sum test, after Bonferroni correction for four possible comparisons).

These results therefore provide experimental support for all four plasticity rules predicted by our proposed mechanism of stochastic gradient descent. We argue in the Discussion that the apparent contradiction of these results with the literature is not unexpected if the likely effects of our altered conditions are considered in the light of known mechanisms of potentiation and depression.

Extraction of the change of error

Above we have provided experimental evidence in support of the counterintuitive synaptic plasticity rules predicted by our proposed learning mechanism. In that mechanism, following a perturbation complex spike, the sign of plasticity is determined by the absence or presence of a follow-up error complex spike that signals whether the movement error increased (spike present) or decreased (spike absent). We now return to the outstanding problem of finding a mechanism able to extract this *change* of error, $\delta\mathcal{E}$.

Several roughly equivalent schemes have been proposed (*Williams, 1992*), including subtraction of the average error (*Barto et al., 1983*) and decorrelation (*Dean and Porrill, 2014*), a specialisation of the covariance rule (*Sejnowski, 1977*). However, in general, these suggestions have not extended
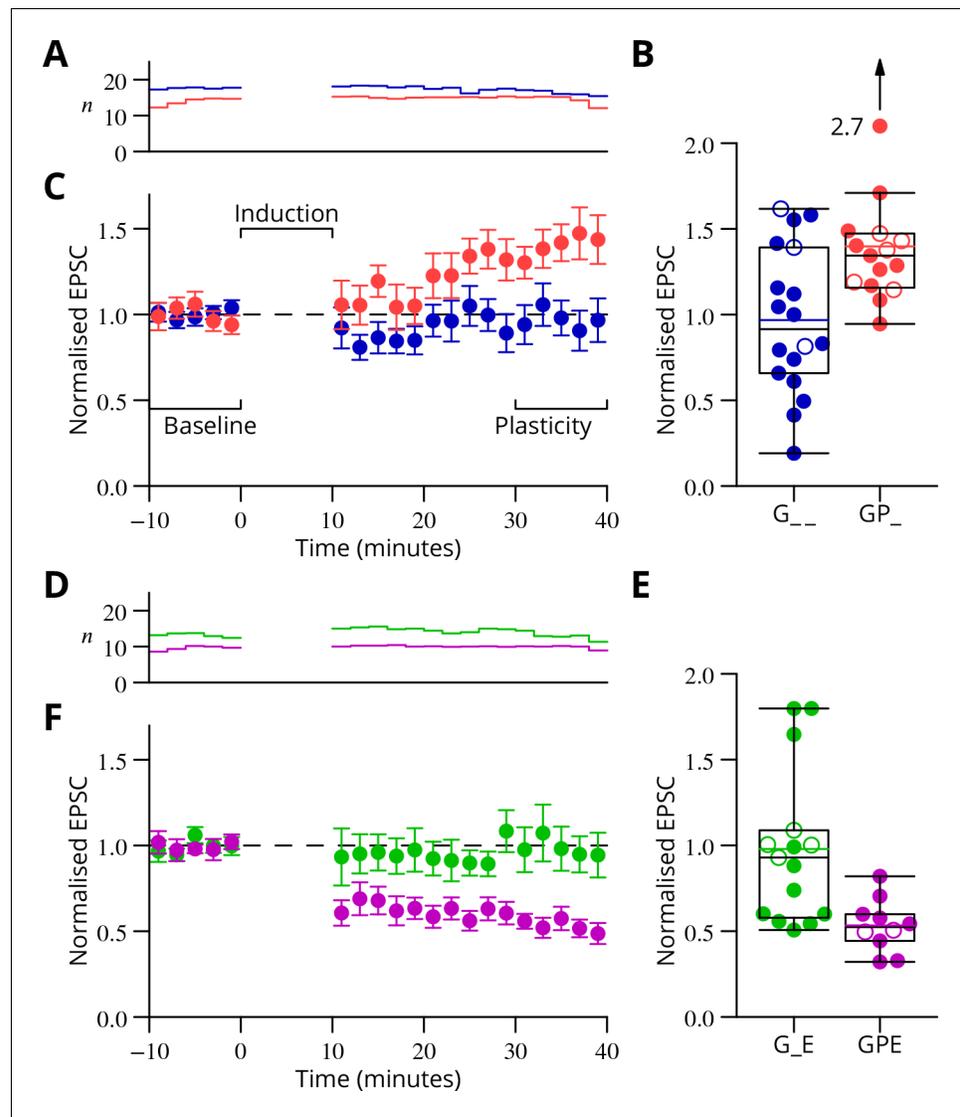
Figure 6. Time course and amplitude of plasticity. (**A**) number, (**B**) box-and-whisker plots of individual plasticity ratios (*coloured lines* represent the means, *open symbols* represent cells with failures of climbing fibre stimulation; see Materials and methods) and (**C**) time course of the mean EPSC amplitude for GP_ (red) and G_ _ (blue) protocols of *Figure 4*, normalised to the pre-induction amplitude. Averages every 2 min, mean ± sem. Non-integer n arise because the numbers of responses averaged were normalised by those expected in two minutes, but some responses were excluded (see Materials and methods) and some recordings did not extend to the extremities of the bins. Induction lasted for 10 min starting at time 0. (**D, E**) and (**F**) similar plots for the GPE (*magenta*) and G_E (*green*) protocols of *Figure 5*.

DOI: https://doi.org/10.7554/eLife.31599.007

The following figure supplement is available for figure 6:

Figure supplement 1. In complementary experiments, we examined the outcome of applying a burst of six climbing fibre stimuli simultaneously with the granule cell burst.

DOI: https://doi.org/10.7554/eLife.31599.008

to detailed cellular implementations. In order to restrict our implementation to biologically plausible mechanisms we selected a method that involves subtracting the average error from the trial-to-trial error (*Barto et al., 1983*; *Doya and Sejnowski, 1988*). The residual of the subtraction is then simply the variation of the error $\delta\mathcal{E}$ as desired.

Table 1. Group data and statistical tests for plasticity outcomes.

In the upper half of the table, the ratios of EPSC amplitudes after/before induction are described and compared with a null hypothesis of no change (ratio = 1). The GP_ and GPE protocols both induced changes, while the control protocols (G_ _, G_E) did not. The bottom half of the table analyses differences of those ratios between protocols. The 95% confidence intervals (c.i.) were calculated using bootstrap methods, while the p-values were calculated using a two-tailed Wilcoxon rank sum test. The p-values marked with an asterisk have been corrected for a two-stage analysis by a factor of 2 (Materials and methods).

Comparison	Mean	95 % c.i.		p	n
GP_	1.40	1.26,	1.72	0.0001	15
GPE	0.53	0.45,	0.63	0.002	10
G_ _	0.97	0.78,	1.16	0.77	18
G_E	0.98	0.79,	1.24	0.60	15
GP_ vs G_ _	0.43	0.19,	0.74	0.021*	
GPE vs G_E	−0.44	−0.72,	−0.24	0.0018*	
GP_ vs G_E	0.42	0.14,	0.73	0.01*	
GPE vs G_ _	−0.43	−0.65,	−0.22	0.008*	
G_ _ vs G_E	0.01	−0.26,	0.32	0.93	

DOI: https://doi.org/10.7554/eLife.31599.009

As mechanism for this subtraction, we propose that the excitatory synaptic drive to the inferior olive is on average balanced by input from the GABAergic nucleo-olivary neurones. A diagram is shown in *Figure 7* to illustrate how this might work in the context of repetitions of a single
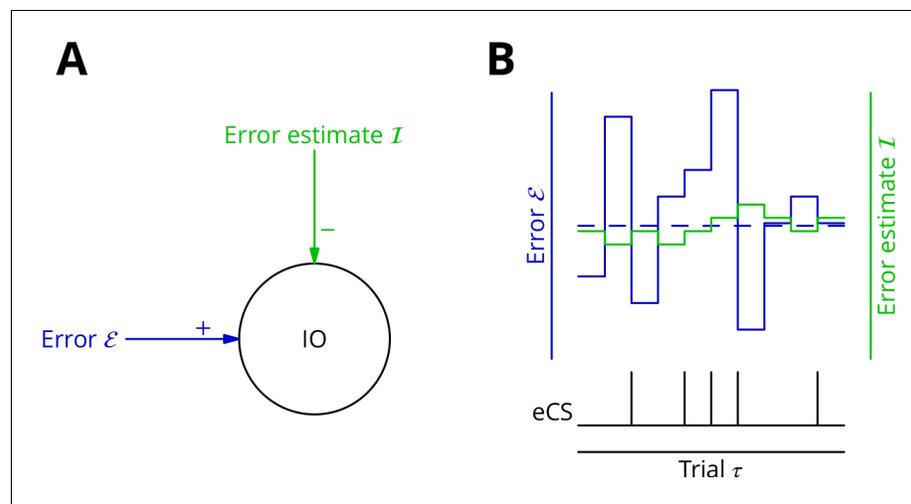
Figure 7. Adaptive tracking to cancel the mean error input to the inferior olive. (**A**) The olive is assumed to receive an excitatory signal representing movement error $\mathcal{E}$ and an inhibitory input $\mathcal{I}$ from the nucleo-olivary neurones of the cerebellar nuclei. (**B**) The inputs to the inferior olive are represented in discrete time (τ)—each bar can be taken to represent a discrete movement realisation. The error (*blue*) varies about its average (*dashed blue line*) because perturbation complex spikes influence the movement and associated error randomly. The strength of the inhibition is shown by the *green* trace. When the excitatory error input exceeds the inhibition, an error complex spike is emitted (*bottom black trace*) and the inhibition is strengthened by plasticity, either directly or indirectly. In the converse situation and in the consequent absence of an error complex spike, the inhibition is weakened. In this way, the inhibition tracks the average error and the emission of an error complex spike signals an error exceeding the estimated average. Note that spontaneous perturbation complex spikes are omitted from this diagram.
DOI: https://doi.org/10.7554/eLife.31599.010

movement (we extend the mechanism to multiple interleaved movements below). Briefly, a feedback plasticity reinforces the inhibition whenever it is too weak to prevent an error complex spike from being emitted. When the inhibition is strong enough to prevent an error complex spike, the inhibition is weakened. If the variations of the strength of the inhibition are sufficiently small and the error does not change rapidly, the level of inhibition will attain a good approximation of the average error. Indeed, this mechanism can be viewed as maintaining an estimate of the movement error. However, the error still varies about its mean on a trial-to-trial basis because of the random perturbations that influence the movement and therefore the error. In consequence, error complex spikes are emitted when the error exceeds the (estimated) average; this occurs when the perturbation increases the error. This mechanism enables extraction of the sign of $\delta\mathcal{E}$ in the context of a single movement realisation. In support of such a mechanism, there is evidence that inhibition in the olive builds up during learning and reduces the probability of complex spikes (*Kim et al., 1998*).

More than one plasticity mechanism could produce the desired cancellation of excitatory drive to the inferior olive. We outline two possibilities here, but it will be necessary in the implementation below to make a concrete if somewhat arbitrary choice; we shall make it on the basis of the available, circumstantial evidence.

The first possible mechanism would involve plasticity of the inhibitory synapses made by nucleo-olivary neurones in the inferior olive (*Figure 1A,B*). Perturbation and error complex spikes would be distinguished in an appropriate plasticity rule by the presence of excitatory synaptic input to the olive. This would offer a simple implementation, since plastic and cancelled inputs would be at neighbouring synapses (*De Zeeuw et al., 1998*); information about olivary spikes would also be directly available. However, the lack of published evidence and our own unsuccessful exploratory experiments led us to consider an alternative plasticity locus.

A second possible implementation for cancelling the average error signal would make the mossy fibre to nucleo-olivary neurone synapses plastic. The presence of an error complex spike would need to potentiate these inputs, thereby increasing inhibitory drive to the olive and tending to reduce the likelihood of future error complex spikes being emitted. Inversely, the absence of the error complex spike should depress the same synapses. Movement specificity could be conferred by applying the plasticity only to active mossy fibres, the patterns of which would differ between movements. This would enable movement-specific cancellation as long as the overlap between mossy fibre patterns was not too great.

How would information about the presence or absence of the error complex spike be supplied to the nucleo-olivary neurones? A direct connection between climbing fibre collaterals and nucleo-olivary neurones exists (*De Zeeuw et al., 1997*), but recordings of cerebellar neurones following stimulation of the olive suggest that this input is not strong, probably eliciting no more than a single spike per activation (*Bengtsson et al., 2011*). The function of this apparently weak input is unknown.

An alternative route to the cerebellar nuclear neurones for information about the error complex spike is via the Purkinje cells. Climbing fibres excite Purkinje cells which in turn inhibit cerebellar nuclear neurones, in which a strong inhibition can cause a distinctive rebound of firing (*Llinás and Mühlethaler, 1988*). It has been reported that peripheral stimulation of the climbing fibre receptive field, which might be expected to trigger the emission of error complex spikes, causes large IPSPs and an excitatory rebound in cerebellar nuclear neurones (*Bengtsson et al., 2011*). These synaptically induced climbing fibre–related inputs were stronger than spontaneously occurring IPSPs. In our conceptual framework, this could be interpreted as indicating that error complex spikes are stronger and/or arise in a greater number of olivary neurones than perturbation complex spikes. The two types of complex spike would therefore be distinguishable, at least in the cerebellar nuclei.

Plasticity of active mossy fibre inputs to cerebellar nuclear neurones has been reported which follows a rule similar to that our implementation requires. Thus, mossy fibres that burst before a hyperpolarisation (possibly the result of an error complex spike) that triggers a rebound have their inputs potentiated (*Pugh and Raman, 2008*), while mossy fibres that burst without a succeeding hyperpolarisation and rebound are depressed (*Zhang and Linden, 2006*). It should be noted, however, that this plasticity was studied at the input to projection neurones and not at that to the nucleo-olivary neurones. Nevertheless, the existence of highly suitable plasticity rules in a cell type closely related to the nucleo-olivary neurones encouraged us to choose the cerebellar nuclei as the site of the plasticity that leads to cancellation of the excitatory input to the olive.

We now consider how synaptic integration in the olive leads to emission or not of error complex spikes. The nucleo-olivary synapses (in most olivary nuclei) display a remarkable degree of delayed and long-lasting release (*Best and Regehr, 2009*), suggesting that inhibition would build up during a command and thus be able to oppose the excitatory inputs signalling movement errors that appear some time after the command is transmitted. The error complex spike would therefore be produced (or not) after the command. On this basis, we shall make the simplifying assumption that the cerebellum generates a relatively brief motor control output or 'command', of the order of 100 ms or less and a single error calculation is performed after the end of that command. As for the saccade example previously mentioned, many movements can only be evaluated after completion. In effect, this corresponds to an offline learning rule.

Simulations

Above we outlined a mechanism for extracting the error change $\delta\mathcal{E}$; it is based on adapting the inhibitory input to the inferior olive to cancel the average excitatory error input in a movement-specific manner. To verify that this mechanism could operate successfully in conjunction with the scheme for cortical plasticity already described, we turned to simulation.

A reduced model of a cerebellar microzone was developed and is described in detail in the Materials and methods. In overview, mossy fibre input patterns drove Purkinje and cerebellar nuclear neurones during commands composed of 10 discrete time bins. Purkinje cell activity was perturbed by randomly occurring complex spikes, which each increased the firing in a single time bin. The learning task was to adjust the output patterns of the nuclear projection neurones to randomly chosen targets. Cancellation of the average error was implemented by plasticity at the mossy fibre to nucleo-olivary neurone synapse while modifications of the mossy fibre pathway input to Purkinje cells reflected the rules for stochastic gradient descent described earlier. Synaptic weights were updated offline after each command. The global error was the sum of absolute differences between the projection neurone outputs and their target values. Error complex spikes were broadcast to all Purkinje cells when the error exceeded the integral of inhibitory input to the olive during the movement. There were thus 400 (40 projection neurones × 10 time bins) variables to optimise using a single error value.

The progress of the simulation is illustrated in *Figure 8*, in which two different movements were successfully optimised in parallel; only one is shown. The error can be seen to decrease throughout the simulation, indicating the progressive improvement of the learnt command. The effect of learning on the difference between the output and the target values can be seen in *Figure 8C and D*. The initial match is very poor because of random initial weight assignments, but it has become almost exact by the end of the simulation.

The optimisation only proceeds once the inhibitory input to the olive has accurately subtracted the average error. Thus, in *Figure 8B* it can be seen that the initial values of the inhibitory and excitatory (error) inputs to the olive differed. The inhibition tends towards the error. Until the two match, the overall error shows no systematic improvement. This confirms the need for accurate subtraction of the mean error to enable extraction of the error changes necessary to descend the error gradient. This simulation supports the feasibility of our proposed cerebellar implementation of stochastic gradient descent.

Algorithm convergence and capacity

The simulations above provide a proof of concept for the proposed mechanism of cerebellar learning. Nonetheless, even for the relatively simple network model in the simulations, it is by no means obvious to determine the regions of parameter space in which the model converges to the desired outputs. It is also difficult to analyse the algorithm's performance compared to other more classical ones. To address these issues, we abstract the core mechanism of stochastic gradient descent with estimated global errors in order to understand better the algorithm dynamics and to highlight the role of four key parameters. Analysis of this mechanism shows that this algorithm, even in this very reduced form, exhibits a variety of dynamical regimes, which we characterise. We then show how the key parameters and the different dynamical learning regimes directly appear in an analog perceptron description of the type considered in the previous section. We find that the algorithm's storage capacity is similar to the optimal capacity of an analog perceptron.
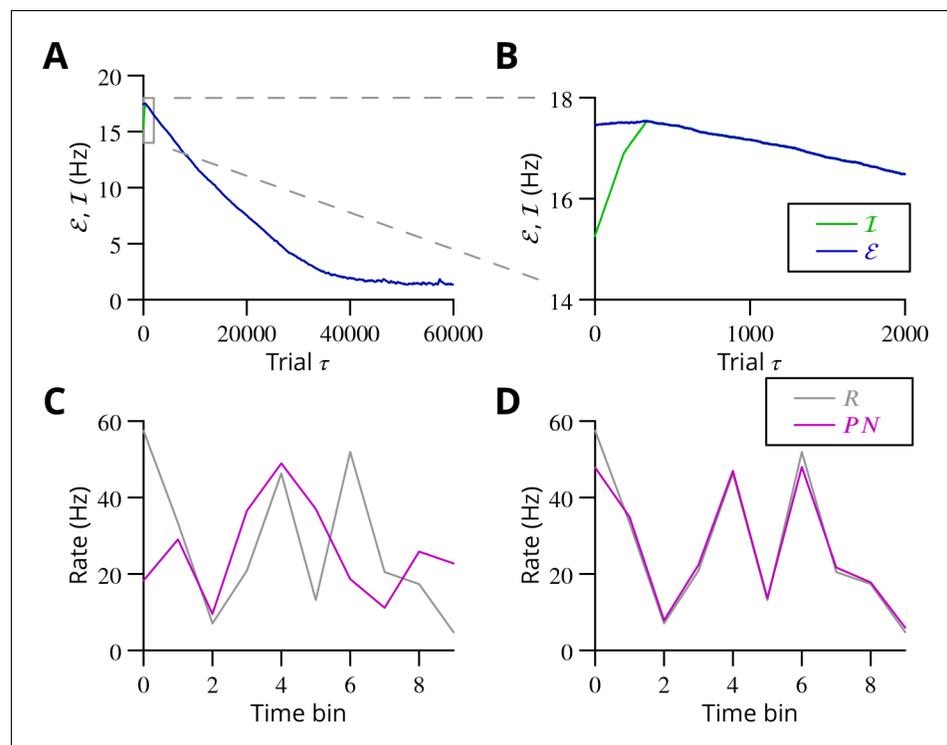
Figure 8. Simulated cerebellar learning by stochastic gradient descent with estimated global errors. The total error ($\mathcal{E}$, *blue*) at the cerebellar nuclear output and the cancelling inhibition ($\mathcal{I}$, *green*) reaching the inferior olive are plotted as a function of trial number (τ) in (**A** and **B**) for one of two interleaved patterns learnt in parallel. An approximately 10-fold reduction of error was obtained. It can be seen in A that the cancelling inhibition follows the error very closely over most of the learning time course. However, the zoom in B shows that there is no systematic reduction in error until the inhibition accurately cancels the mean error. (**C**) Initial firing profile of a typical cerebellar nuclear projection neurone (*PN*, *magenta*); the simulation represented 10 time bins with initially random frequency values per neurone, with a mean of 30 Hz. The target firing profile for the same neurone (*R*, *grey*) is also plotted. (**D**) At the end of the simulation, the firing profile closely matched the target.
DOI: https://doi.org/10.7554/eLife.31599.011

Reduced model

In order to explore more exhaustively the convergence of the learning algorithm, we considered a reduced version with a single principal cell that nevertheless captures its essence (**Figure 9A**). A detailed analysis of this circuit is presented in the Appendix 1, which we summarise and illustrate here.

We focus on the rate P of a single P-cell (principal cell, in essence a combination of a Purkinje cell and a nuclear projection neurone), without considering how the firing is driven by synaptic inputs. The other variable of the model is the strength $\mathcal{J}$ of the plastic excitatory input from mossy fibres to nucleo-olivary neurones **Figure 9**. The learning task is to bring the firing rate P of the P-cell to a desired target rate R, guided by an estimation of the current error relying on $\mathcal{J}$, which is refined concurrently with P from trial to trial.

The error $\mathcal{E}$ is determined by the difference between P and R; we choose the absolute difference $\mathcal{E} = |P - R|$. In the presence of a perturbation A, occurring with probability ρ, the rate becomes $P + A$ and so $\mathcal{E} = |P + A - R|$. The current estimate of the error is measured by the strength $\mathcal{I}$ of the inhibition on the olivary neurone associated with the P-cell. It assumes the value $\mathcal{I} = [\mathcal{J} - qP]_+$ (the brackets $[X]_+$ denoting the rectified value of X, $[X]_+ = X$ if $X \geq 0$ and 0 otherwise); q represents the strength of P-cell (in reality Purkinje cell) synapses on nucleo-olivary neurones. The inhibition of the inferior olive (IO) arises from the discharge of the nucleo-olivary neurones (NO) induced by the mossy fibre inputs, the $\mathcal{J}$ component of $\mathcal{I}$. The NOs are themselves inhibited by the P-cell which is accounted for by the $-qP$ component of $\mathcal{I}$ (**Figure 9A**). Note that in the presence of a perturbation
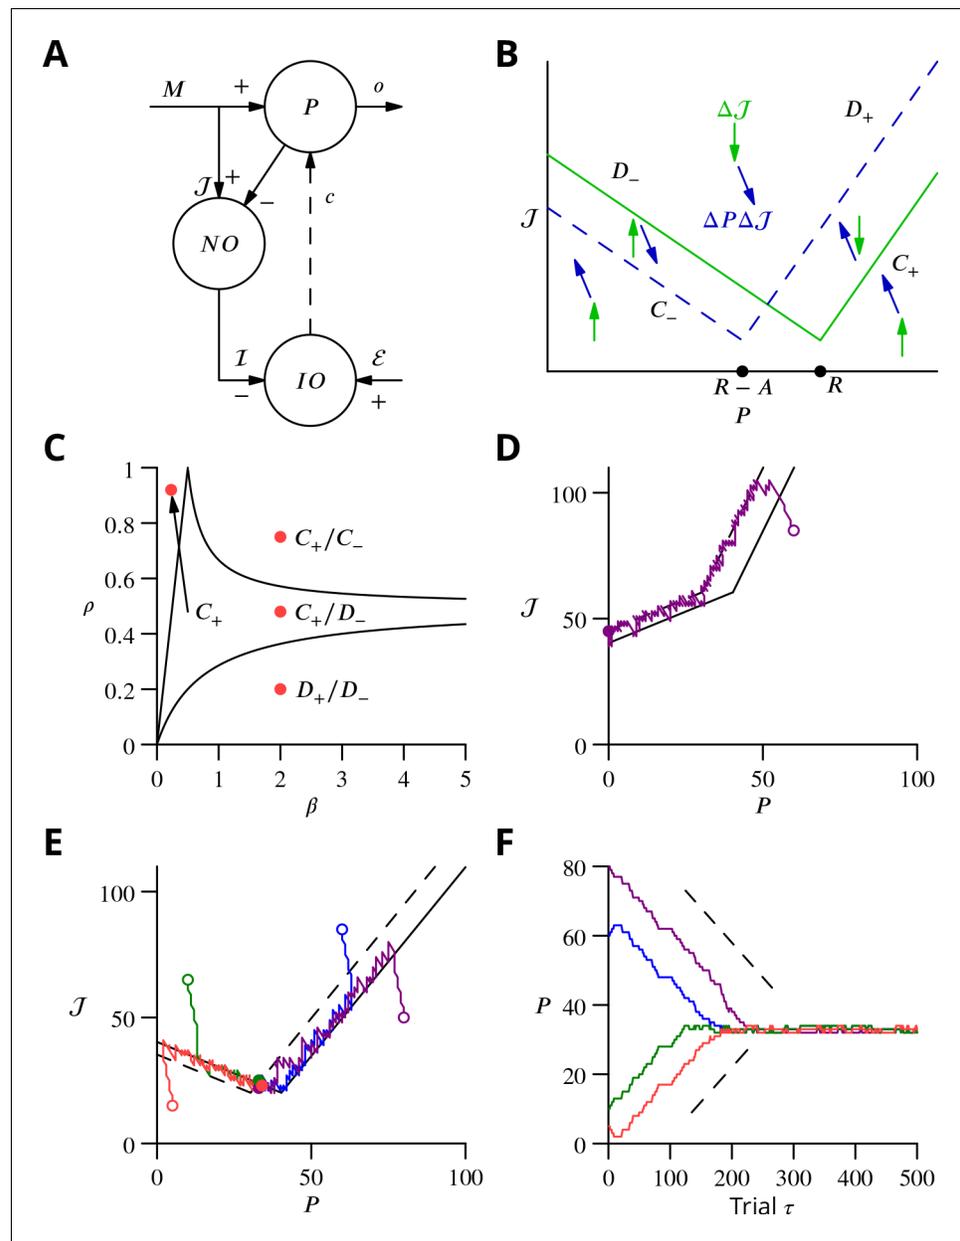
Figure 9. Single-cell convergence dynamics in a reduced version of the algorithm. (**A**) Simplified circuitry implementing stochastic gradient descent with estimated global errors. Combined Purkinje cells/projection neurones (P) provide the output o. These cells receive an excitatory plastic input from mossy fibres (M). Mossy fibres also convey a plastic excitatory input $\mathcal{J}$ to the nucleo-olivary neurones (NO), which receive inhibitory inputs from the Purkinje cells and supply the inhibitory input $\mathcal{I} = [\mathcal{J} - qP]_+$ to the inferior olive (IO). The inferior olive receives an excitatory error input $\mathcal{E}$. The olivary neurones emit spikes that are transmitted to the P-cells via the climbing fibre c. (**B**) Effects of plasticity on the simplified system in the plane defined by the P-cell rate P (with optimum R, perturbation A) and excitatory drive to the nucleo-olivary neurones $\mathcal{J}$. Plastic updates of type $\Delta P \Delta \mathcal{J}$ (*blue arrows*) and $\Delta \mathcal{J}$ (*green arrows*) are shown. The updates change sign on the line C (*dashed blue*) and D (*solid green*), respectively. Lines $C_\pm$ and $D_\pm$ delimit the 'convergence corridors'. The diagram is drawn for the case $q<1$, in which A, the perturbation of P is larger than the perturbation qA of $\mathcal{I}$. (**C**) Parameters β and ρ determine along which lines the system converges to R. The *red dot* in the D_+/D_- region shows the parameter values used in panels E and F and also corresponds to the effective parameters of the perceptron simulation of ***Figure 10***. The other *red dots* show the parameters used in the learning examples displayed in ***Appendix 1—figure 1***. (**D**) When $q>1$, C_+ and D_- do not cross and learning does not converge to the desired rate. After going down the C_+/D_+ corridor, the point $(P, \mathcal{J})$ continues along the C_-/D_- corridor without stopping close to the desired target rate R. *Figure 9 continued on next page*

Figure 9 continued

Open circle, start point; *filled circle*, end point. (**E**) Dynamics in the $(P, \mathcal{J})$ plane for $\rho = 0.2$ and $\beta = 2$
($A = 10, \Delta P = 1, \Delta \mathcal{J} = 2, q = 0.5$). Trajectories (*coloured lines*) start from different initial conditions (*open circles*).
Open circles, start points; *filled circles*, end points. (**F**) Time courses of P as a function of trial τ for the trajectories
in C (*same colours*). *Dashed lines*: predicted rate of convergence (Appendix 1). All trajectories end up fluctuating
around $P = R - A(q + 1)/2$, which is $R - 7.5$ with the chosen parameters. *Open circles*, start points; *filled circles*,
end points.

of the P-cell, the estimate of the error is itself modified to $[\mathcal{J} - qP - qA]_+$ to account for the
decrease (for $q > 0$) of the NO discharge arising from the P-cell firing rate increase.

The system learning dynamics can be analysed in the P-$\mathcal{J}$ plane (*Figure 9B*). After each 'move-
ment' realisation, the two system variables, P and $\mathcal{J}$ are displaced within the P-$\mathcal{J}$ plane by plasticity
according to the algorithm. In the presence of a perturbation, P is decreased (a leftwards displace-
ment) if $\mathcal{E} > [\mathcal{J} - qP - qA]_+$ and otherwise increased. However, P remains unchanged if there is no
perturbation. Conversely, $\mathcal{J}$ is always updated irrespective of the presence or absence of a perturba-
tion. If $\mathcal{E} > [\mathcal{J} - qP - qA]_+$ in the presence of a perturbation or $\mathcal{E} > \mathcal{J} - qP$ in its absence, $\mathcal{J}$ is increased
(upwards displacement) and it is decreased otherwise.

Updates can therefore be described as $\Delta P \Delta \mathcal{J}$ or $\Delta \mathcal{J}$. The resultant directions vary between the
different regions of the $(P, \mathcal{J})$-plane (*Figure 9B*) delimited by the two borders C and D, defined as
$\mathcal{J} - qP - qA = |P - R + A|$ and $\mathcal{J} - qP = |P - R|$, respectively. The two half-lines bordering each sec-
tor are denoted by a plus or minus index according to their slopes.

Stochastic gradient descent is conveniently analysed with a phase-plane description, by following
the values $(P, \mathcal{J})$ from one update to the next. The dynamics randomly alternate between updates
of type $\Delta P \Delta \mathcal{J}$ and $\Delta \mathcal{J}$ which leads $(P, \mathcal{J})$ to follow a biased random walk in the $(P, \mathcal{J})$ plane. Mathe-
matical analysis (see Appendix 1) shows that the dynamics proceed in three successive phases, as
observed in the simulations of *Figure 8*. First, the pair of values $(P, \mathcal{J})$ drifts from the initial condition
towards one of the two 'corridors' between the lines C_+ and D_+ or between C_- and D_- (see
Figure 9B). This first phase leads the estimated error $\mathcal{I}$ to approximate closely the error $\mathcal{E}$, as seen
in the full network simulation (*Figure 8*).

When a corridor is reached, in a second phase, $(P, \mathcal{J})$ follows a stochastic walk in the corridor
with, under suitable conditions, a mean linear decrease of the error in time with bounded fluctua-
tions. The precise line followed and the mean rate of error decrease depends on the initial condi-
tions and two parameters: ρ the probability of a perturbation occurring and $\beta = \Delta \mathcal{J} / \Delta P$, as indicated
in *Figure 9C*. Typical trajectories and time courses for specific values of ρ and β and different initial
conditions are shown in *Figure 9D,E,F* as well as in *Appendix 1—figure 1*.

Error decrease in this second phase requires certain restrictions upon the four key parameters—ρ,
β, ΔP, $\Delta \mathcal{J}$. These restrictions are that $0 < \rho < 1$, which ensures that updates are not restricted to a line,
and $\Delta \mathcal{J} < (1 - q)A$ as well as $\Delta P + \Delta \mathcal{J}/(1 + q) < A$, which are sufficient to ensure that a trajectory will
eventually enter a convergence corridor (large updates might always jump over the corridor).

In a final phase, $(P, \mathcal{J})$ fluctuates around the intersection of C_+ and D_-, when it exists.

C_+ and D_- do not cross for $q > 1$, when the perturbation qA of the estimated error $\mathcal{I}$ is larger than
A the perturbation of the P-cell discharge. In this case, P does not stop close to its target value, as
illustrated in *Figure 9D*. For $q < 1$, C_+ and D_- intersect at $(R - A(q + 1)/2, qR + (1 - q^2)A/2)$. The final
error in the discharge rate fluctuates around $A(q + 1)/2$. Namely, the mean final error grows from
$A/2$ when $q = 0$ (vanishing inhibition of the nucleo-olivary discharge by the P-cell) to A when $q = 1$
(maximal admissible inhibition of the nucleo-olivary discharge by the P-cell). The failed convergence
for $q \geq 1$ could be expected, since in this scenario the perturbation corrupts the error estimate more
than it influences the global error due to the perturbed P-cell firing rate.

Analog perceptron

In order to extend the above analysis to include multiple synaptic inputs and to explore the algo-
rithm's capacity to learn multiple movements, we considered an analog perceptron using the algo-
rithm of stochastic gradient descent with estimated global errors. This allowed us to investigate the
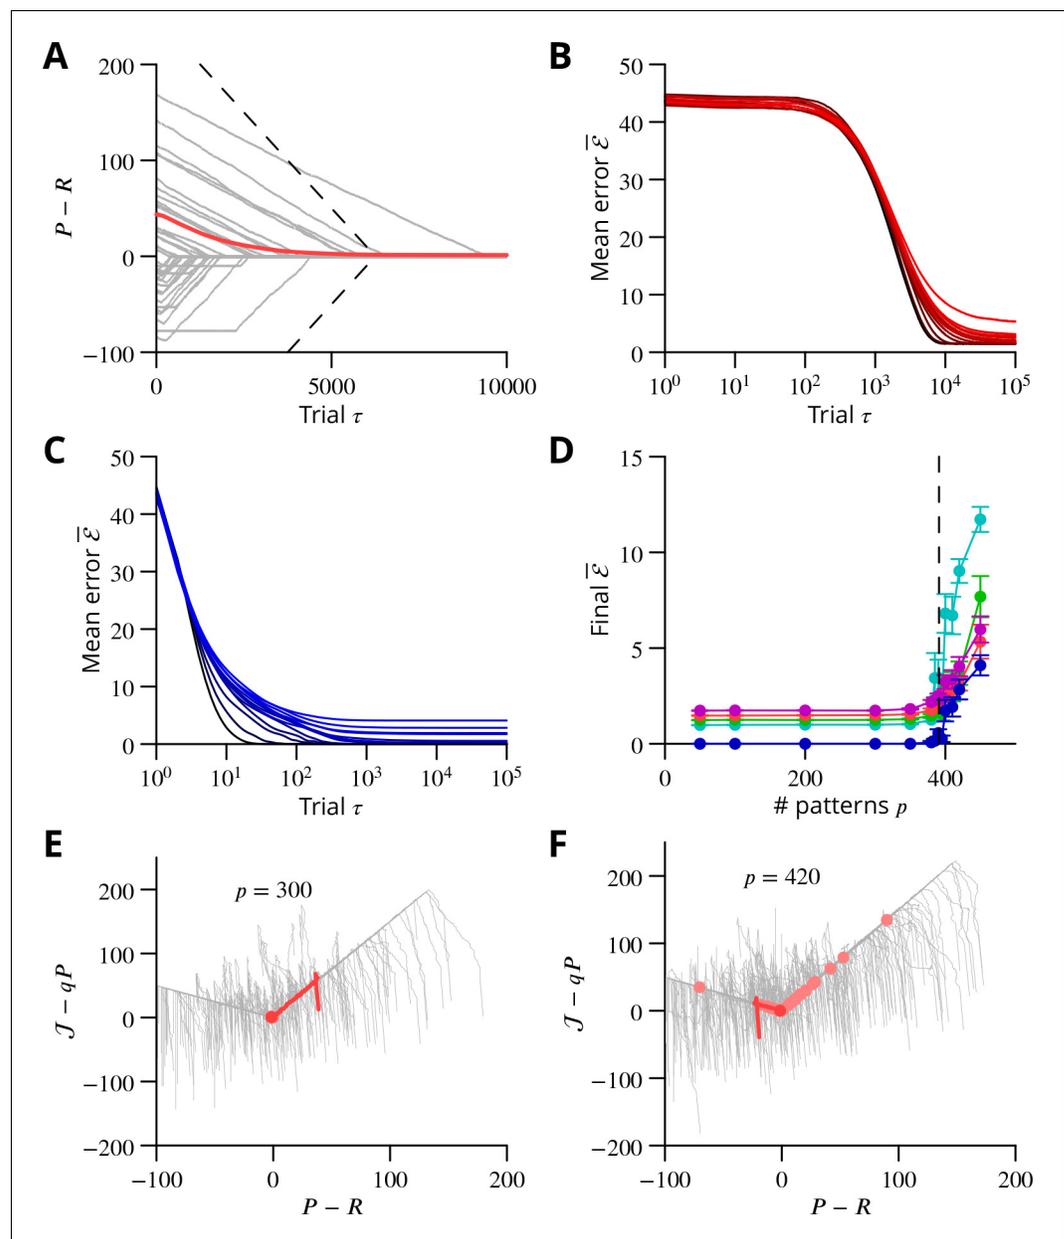
Figure 10. Convergence and capacity for an analog perceptron. (**A**) Convergence for a single learned pattern for stochastic gradient descent with estimated global errors (SGDEGE). Different lines (*thin grey lines*) correspond to distinct simulations. The linear decrease of the error ($P - R$) predicted from the simplified model without zero-weight synapses is also shown (*dashed lines*). When learning increases the rate towards its target, the predicted convergence agrees well with the simulation. When learning decreases the rate towards its target, the predicted convergence is larger than observed because a fraction of synapses have zero weight. (**B**) Mean error vs. number of trials (always per pattern) for different numbers of patterns for SGDEGE with $q = 0.5$. *Colours from black to red*: 50, 100, 200, 300, 350, 380, 385, 390, 395, 400, 410, 420, 450. (**C**) Mean error vs. number of trials when learning using the delta rule. The final error is zero when the number of patterns is below capacity (predicted to be 391 for these parameters), up to finite size effects. *Colours from black to blue*: same numbers of patterns as listed for B. (**D**) Mean error after 1×10^5 trials for the delta rule (*blue*) and for the SGDEGE with $q = 0$ (*cyan*), 0.25 (*green*), 0.5 (*red*), 0.75 (*magenta*), as a function of the number of patterns p. The mean error diverges from its minimal value close to the theoretical capacity (*dashed line*) for both learning algorithms. (**E**) Dynamics of pattern learning in the $P - \mathcal{J}$ or more precisely $(P - R) - (\mathcal{J} - qP)$ plane below maximal capacity ($p = 300$). The error ($P - R$) corresponding to each pattern (*grey*) is reduced as it moves down its convergence corridor. The trajectory for a single specimen pattern is highlighted in *red*, while the endpoints of all trajectories are indicated by *salmon filled*

Figure 10 continued on next page

storage capacity attainable with the algorithm and to compare it to the theoretical maximal capacity for an analog perceptron (*Clopath and Brunel, 2013*).

The architecture was again that of *Figure 9A*; the details of the methods can be found in Appendix 1 and we summarise the conclusions here, with reference to *Figure 10*. The simulation included 1000 inputs (mossy fibres) to a P-cell; note that much larger numbers of inputs pertain in vivo (*Napper and Harvey, 1988*).

As found in the reduced model above, convergence of the algorithm requires the perturbation of the error estimate to be smaller than the effect of the P-cell perturbation on the true error, namely $q<1$. When this condition holds, the rate of learning a single pattern and the final error can be related to that predicted from the simplified analysis above (*Figure 10A* and in Appendix 1).

Learning using perturbations is slower than when using the full error information (i.e., the 'delta rule' (*Dayan and Abbott, 2001*); compare *Figure 10B and C*), for all numbers of patterns. The difference can be attributed to the different knowledge requirement of the two algorithms. When the error magnitude is known, which requires knowledge of the desired endpoint, large updates can be made for large errors, as done in the delta rule. In contrast, the proposed SGDEGE algorithm does not require knowledge of the error absolute magnitude or sign and thus proceeds with constant updating steps. Extensions of the SGDEGE algorithm that incorporate adaptive weight updates and could accelerate learning are mentioned in the Discussion.

The precision of learning is limited by the on-going perturbations to $A(1 + q)/2$ (e.g. $A/2$ for $q = 0$ or, $1.5A$ for $q = 0.5$ as used in *Figure 10B*), but the final average error increases beyond this floor only very close to the theoretical capacity computed for the input-output association statistics (*Figure 10D*). The contribution of interfering patterns to the slowed learning is similar for the two algorithms. The behaviour of individual learning trajectories below maximal capacity is shown in *Figure 10E and F*.

These analyses establish the convergence mechanisms of stochastic gradient descent with estimated global errors and show that it can attain the maximum theoretical storage capacity up to a non-zero final error resulting from the perturbation. Learning is slower than when the full error information is used in the delta rule, but we argue that the availability of that information would not be biologically plausible for most complex movements.

Discussion

A cellular implementation of stochastic gradient descent

Analysis of the requirements and constraints for a general cerebellar learning algorithm highlighted the fact that the current consensus Marr-Albus-Ito model is only capable of learning simple reflex movements. Optimisation of complex, arbitrary movements, of which organisms are certainly capable and to which the cerebellum is widely believed to contribute, would require a different algorithm. We therefore sought to identify within the cerebellar system an implementation of stochastic gradient descent. This should comprise several elements: a source of perturbations, a mechanism for extracting the change of error, and a plasticity rule incorporating this information. We identified a strong constraint on any implementation, requiring each calculation to be made in the context of a single movement realisation. This arises from the potentially arbitrary sequencing of movements with different optima. We also sought a mechanism that only makes use of plausible cellular calculations: summation of excitation and inhibition in the presence of a threshold.

We suggest that the perturbation is provided by the complex spike, which has suitable properties: spontaneous irregular activity, an unambiguous sign during the action potential burst, salience

at a cellular and network level, and the ability to influence synaptic plasticity. This choice of perturbation largely determines the predicted cerebellar cortical plasticity rules: only granule cell inputs active at the same time as a perturbation complex spike undergo plasticity, whose sign is determined by the absence (LTP) or presence (LTD) of a succeeding error complex spike. We have provided evidence that the synaptic plasticity rules do operate as predicted, in vitro under conditions designed to be more physiological than is customary.

An additional plasticity mechanism seems to be required to read off the change of error. The general mechanism we propose involves subtraction of the average error to expose the random variations caused by the perturbations of the movement. The subtraction results from adaptive tracking of the excitatory input to the olivary neurones by the inhibitory input from the nucleo-olivary neurones of the cerebellar nuclei. We chose to place the plasticity at the mossy fibre–nucleo-olivary neurone synapse, mostly because of the existence of suitable plasticity rules at the mossy fibre synapse onto the neighbouring projection neurones. However, plasticity in the olive at the nucleo-olivary input would probably be functionally equivalent and we do not intend to rule out this or alternative sites of the error-cancelling plasticity.

By simulating a simplified cerebellar network implementing this mechanism, we established the ability of our proposed mechanism to learn multiple arbitrary outputs, optimising 400 variables per movement with a single error value. More formal analysis of a simplified version of stochastic gradient descent with estimated global errors established convergence of the algorithm and allowed us to estimate its storage capacity.

Implications for studies of synaptic plasticity

The plasticity rules for parallel fibre–Purkinje cell synapses predicted by our algorithm appeared to be incompatible with the well-established consensus. However, we show that under different, arguably more physiological, conditions, we were able to provide support for the four predicted outcomes.

We made several changes to the experimental conditions, only one of which is specific to the cerebellum. Thus, keeping synaptic inhibition intact has long been recognised as being of potential importance, with debates regarding its role in hippocampal LTP dating back decades (*Wigström and Gustafsson, 1983a*; *Wigström and Gustafsson, 1983b*; *Arima-Yoshida et al., 2011*). Very recent work also highlights the importance of inhibition in the induction of cerebellar plasticity (*Rowan et al., 2018*; *Suvrathan and Raymond, 2018*).

We also made use of a lower extracellular calcium concentration than those almost universally employed in studies of plasticity in vitro. In vivo measurements of the extracellular calcium concentration suggest that it does not exceed 1.5 mM in rodents, yet most studies use at least 2 mM. A 25% alteration of calcium concentration could plausibly change plasticity outcomes, given the numerous nonlinear calcium-dependent processes involved in synaptic transmission and plasticity (*Nevian and Sakmann, 2006*; *Graupner and Brunel, 2007*).

A major change of conditions we effected was cerebellum-specific. Nearly all studies of granule cell–Purkinje cell plasticity have employed stimulation of parallel fibres in the molecular layer. Such concentrated, synchronised input activity is unlikely to arise physiologically. Instead of this, we stimulated in the granule cell layer, a procedure expected to generate a more spatially dispersed input on the Purkinje cell, presumably leading to minimised dendritic depolarisations. Changing the stimulation method has been reported to prevent induction of LTD using standard protocols (*Marcaggi and Attwell, 2007*).

Although we cannot predict in detail the mechanistic alterations resulting from these changes of conditions, it is nevertheless likely that intracellular calcium concentrations during induction will be reduced, and most of the changes we observed can be interpreted in this light. It has long been suggested that high calcium concentrations during induction lead to LTD, while lower calcium concentrations generate LTP (*Coesmans et al., 2004*); we have recently modelled the induction of this plasticity, incorporating both calcium and nitric oxide signalling (*Bouvier et al., 2016*). Consistently with this viewpoint, a protocol that under standard conditions produce LTD—simultaneous activation of granule cells and climbing fibres—could plausibly produce LTP in the present conditions as a result of reduced intracellular calcium. Analogously, granule cell stimulation that alone produces LTP under standard conditions might elicit no change if calcium signalling were attenuated under our conditions.

Interestingly, LTP resulting from conjunctive granule cell and climbing fibre stimulation has been previously reported, in vitro (*Mathy et al., 2009*; *Suvrathan et al., 2016*) and in vivo (*Wetmore et al., 2014*). In contrast, our results do not fit well with several other studies of plasticity in vivo (*Ito et al., 1982*; *Jörntell and Ekerot, 2002*; *Jörntell and Ekerot, 2003*; *Jörntell and Ekerot, 2011*). However, in these studies quite intense stimulation of parallel and/or climbing fibre inputs was used, which may result in greater depolarisations and calcium entry than usually encountered. This difference could therefore account for the apparent discrepancy with the results we predict and have found in vitro.

It is unlikely that the interval between perturbation and error complex spikes would be fixed from trial to trial and it is certainly expected to vary with sensory modality (vision being slow). Our theoretical framework would therefore predict that a relatively wide range of intervals should be effective in inducing LTD through the GPE protocol. However, this prediction is untested and it remains possible that different intervals (potentially in different cerebellar regions) may lead to different plasticity outcomes, as suggested by the recent work of *Suvrathan et al. (2016)*; this might also contribute to some of the variability of individual plasticity outcomes we observe.

Another open question is whether different relative timings of parallel and climbing fibre activity would result in different plasticity outcomes. In particular, one might hypothesise that parallel fibres active during a pause following a perturbation complex spike might display plasticity of the opposite sign to that reported here for synchrony with the complex spike itself.

In summary, while in vitro studies of plasticity are likely to reveal molecular mechanisms leading to potentiation and depression, the outcomes from given stimulation protocols may be very sensitive to the precise conditions, making it difficult to extrapolate to the in vivo setting, as we have shown here for the cerebellum. Similar arguments could apply to in vitro plasticity studies in other brain regions.

Current evidence regarding stochastic gradient descent

As mentioned in the introductory sections, the general cerebellar learning algorithm we propose here is not necessarily required in situations where movements are simple or constrained, admitting a fixed mapping between errors and corrective action. Furthermore, such movements constitute the near totality of well-studied models of cerebellar learning. Thus, the vestibulo-ocular reflex and saccade adaptation involve eye movements, which are naturally constrained, while the eyeblink is a stereotyped protective reflex. There is therefore a possibility that our mechanism does not operate in the cerebellar regions involved in oculomotor behaviour, even if it does operate elsewhere.

In addition, these ocular behaviours apparently display error functions that are incompatible with our assumptions. In particular, disturbance of a well optimised movement would be expected to increase error. However, it has been reported multiple times that climbing fibre activity can provide directional error information, including *reductions* of climbing fibre activity below baseline (e.g. *Soetedjo et al., 2008*). This argument is not totally conclusive, however. Firstly, we recall that the error is represented by the input to the inferior olive, not its output. It is thus possible that inputs from the nucleo-olivary neurones (or external inhibitory inputs) to the olive also have their activity modified by the disturbance of the movement, causing the reduction of climbing fibre activity. Secondly, what matters for our algorithm is the temporal sequence of perturbation and error complex spikes, but investigation of these second-order statistics of complex spike activity in relation to plasticity has, to our knowledge, not been reported. Similarly, it has been reported that learning and plasticity (LTD) occur in the absence of *modulation* of climbing fibre activity (*Ke et al., 2009*). Although this is difficult to reconcile with either the standard theory or our algorithm, it does not entirely rule out the existence of perturbation-error complex spike pairs that we predict lead to LTD.

Trial-to-trial plasticity correlated with the recent history of complex spikes has been demonstrated in oculomotor adaptation (*Medina and Lisberger, 2008*; *Yang and Lisberger, 2014*). This suggests that one way of testing whether our algorithm operates would be to examine whether the history of complex spike activity can predict future changes in the simple spike firing rate, according to the plasticity rules described above. For instance, two complex spikes occurring at a short interval should cause at the time of the first an increase of simple spike firing in subsequent trials. However, the most complete datasets reported to date involve oculomotor control (*Yang and Lisberger, 2014*; *Catz et al., 2005*; *Soetedjo et al., 2008*; *Ke et al., 2009*), where, as mentioned above, our algorithm may not be necessary.

Beyond the predictions for the plasticity rules at parallel fibre–Purkinje cell synapses tested above, there are a number of aspects of our theory that do fit well with existing observations. The simple existence of spontaneous climbing fibre activity is one. Additional suggestive features concern the evolution of climbing fibre activity during eyeblink conditioning (*Ohmae and Medina, 2015*). Once conditioning has commenced, the probability of complex spikes in response to the unconditioned stimulus decreases, which would be consistent with the build-up of the inhibition cancelling the average error signal in the olive. Furthermore, omission of the unconditioned stimulus then causes a reduction in the probability of complex spikes below the baseline rate, strongly suggesting a specifically timed inhibitory signal has indeed developed at the time of the unconditioned stimulus (*Kim et al., 1998*).

We suggest that the cancellation of average error involves plasticity at mossy fibre–nucleo-olivary neurone synapses. To date no study has reported such plasticity, but the nucleo-olivary neurones have only rarely been studied. Plasticity at the mossy fibre synapses on projection neurones has been studied both in vitro (*Pugh and Raman, 2006*; *Pugh and Raman, 2008*; *Zhang and Linden, 2006*) and in vivo (*Ohyama et al., 2006*), but is not used in our proposed algorithm. Axonal remodelling and synaptogenesis of mossy fibres in the cerebellar nuclei may underlie this plasticity (*Kleim et al., 2002*; *Boele et al., 2013*; *Lee et al., 2015*) and could also contribute to the putative plasticity at mossy fibre synapses on nucleo-olivary neurones.

Finally, our theory of course predicts that perturbation complex spikes perturb ongoing movements. It is well established that climbing fibre activation can elicit movements (*Barmack and Hess, 1980*; *Kimpo et al., 2014*; *Zucca et al., 2016*), but it remains to be determined whether the movements triggered by spontaneous climbing fibre activity are perceptible. *Stone and Lisberger (1986)* reported the absence of complex-spike-triggered eye movements in the context of the vestibulo-ocular reflex. However, it is known that the visual system is very sensitive to retinal slip (*Murakami, 2004*), so it may be necessary to carry out high-resolution measurements and careful averaging to confirm or exclude the existence of perceptible movement perturbations.

Climbing fibre receptive fields and the bicycle problem

If the cerebellum is to contribute to the optimisation of complex movements, its output controlling any given muscle must be adjustable by learning involving multiple error signals. Thus, it may be useful to adjust an arm movement using vestibular error information as part of a balancing movement, but using touch information in catching a ball. However, it is currently unclear to what extent individual Purkinje cells can access different error signals.

There is an extensive literature characterising the modalities and receptive fields of climbing fibres. The great majority of reports are consistent with climbing fibres having fixed, specific modalities or very restricted receptive fields, with neighbouring fibres having similar properties (*Garwicz et al., 1998*; *Jörntell et al., 1996*). Examples would be a climbing fibre driven by retinal slip in a specific direction (*Graf et al., 1988*) or responding only to a small patch of skin (*Garwicz et al., 2002*). These receptive fields are quite stereotyped and have proven to be reliable landmarks in the functional regionalisation of the cerebellum; they are moreover tightly associated with the genetically specified zebrin patterning of the cerebellum (*Schonewille et al., 2006b*; *Mostofi et al., 2010*; *Apps and Hawkes, 2009*).

The apparently extreme specialisation of climbing fibres would limit the ability of the cerebellar circuitry to optimise complex movements. We can illustrate this with a human behaviour: riding a bicycle, which is often taken as an example of a typical cerebellar behaviour. This is an acquired skill for which there is little evolutionary precedent. It is likely to involve learning somewhat arbitrary arm movements in response to vestibular input (it is possible to ride a bike with one's eyes closed). The error signals guiding learning could be vestibular, visual or possibly cutaneous/nociceptive (as a result of a fall), but not necessarily those related to the arm whose movement is learnt. How can such disparate or uncommon but sometimes essential error signals contribute to cerebellar control of the arm? We call this the 'bicycle problem'.

At least two, non-exclusive solutions to this problem can be envisaged (see *Figure 11*). The first we term 'output convergence'; it involves the convergence of multiple cerebellar regions receiving climbing fibres of different modalities onto each specific motor element (for instance a muscle) being controlled. Striking, if partial, evidence for this is found in a study by (*Ruigrok et al., 2008*), who injected the retrograde, trans-synaptic tracer rabies virus into individual muscles. Multiple cerebellar
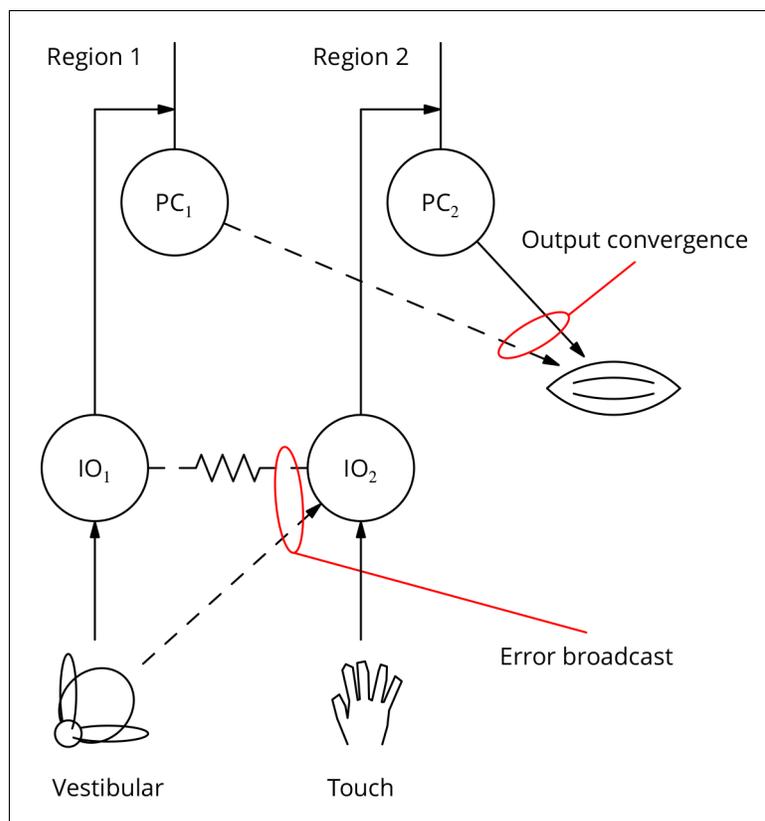
Figure 11. Diagram illustrating two possible solutions to the 'bicycle problem': how to use vestibular error information to guide learning of arm movements to ride a bicycle. In the 'output convergence' solution, the outputs from cerebellar regions receiving different climbing fibre modalities converge onto a motor unit (represented by a muscle in the diagram). In the 'error broadcast' solution, error complex spikes are transmitted beyond their traditional receptive fields, either by divergent synaptic inputs and/or via the strong electrical coupling between inferior olivary neurones.

zones were labelled, showing that they all contribute to the control of those muscles, as posited. What is currently less clear is whether such separate zones receive climbing fibre inputs with different modalities. We note that the output convergence solution to the bicycle problem implies that the synaptic changes in those regions receiving the appropriate error information must outweigh any drift in synaptic weights from those regions deprived of meaningful error information. Adaptation of the learning rate, an algorithmic extension we suggest below, could contribute to ensuring that useful synaptic changes dominate.

We term the second solution to the bicycle problem the 'error broadcast' solution. According to this, error inputs to the olive are broadcast to olivary neurones (and Purkinje cells) outside the traditional receptive field. Although the weight of literature appears to be against this, there are both possible mechanisms and a small amount of supporting data for this suggestion. In terms of mechanism, the well-known electrical coupling of olivary neurones (*Devor and Yarom, 2002*) could recruit cells that do not directly receive suprathreshold synaptic input. This may occur much more frequently in vivo than in the quiescent/anesthetised conditions employed for most studies of climbing fibre receptive fields. Evidence for 'broadcast' of what we would term error complex spikes in vivo has been reported for visual and auditory stimuli (*Mortimer, 1975*; *Ozden et al., 2012*); these stimuli may be correlated with startle responses. Eye blink conditioning using a visual unconditioned stimulus has also been reported (*Rogers et al., 1999*).

The existence of broadcast error complex spikes would provide a mechanism explaining the giant IPSPs in the cerebellar nuclei elicited by peripheral stimulation (*Bengtsson et al., 2011*) and could

also account for the correlation between visual smooth pursuit behaviour and the nature of individual complex spikes (*Yang and Lisberger, 2014*): behaviour could only correlate with a single cell if others were receiving the same input.

Possible extensions to the algorithm

Our implementation of cerebellar stochastic gradient descent and its simulation were purposefully kept as simple as possible, to provide a proof-of-concept with a minimum of assumptions and to simplify parallel analyses. It is likely that parts of the implementation will need to be altered and/or extended as further information becomes available.

Probably the most uncertain element of the implementation is the adaptive site in the cancellation of the average error. We chose to make the mossy fibre–nucleo-olivary neurone synapse plastic, but the plasticity could certainly operate in the olive instead of or in addition to the cerebellar nuclear site. Further studies of synaptic transmission and plasticity in both structures are clearly warranted in this context.

A simplification in our implementation is that it represents brief, discrete commands in what amounts to an offline learning rule. Error complex spikes are only emitted after the command and indeed were not simulated explicitly. This has the great advantage of avoiding the question of whether Purkinje cells that have not received a perturbation complex spike would interpret a broadcast error complex spike as a perturbation. A first remark is to stress the fact that the movement evaluation will often only occur after the movement is complete. Even if an eCS caused a small movement or some additional plasticity, that might be less of an issue outside of a movement context. Additionally, it is possible that cellular mechanisms could exist that would enable the Purkinje cell to distinguish the two types of input and therefore avoid spurious plasticity. The most obvious mechanism would be that, as already hinted at in the literature, error complex spikes are likely to be stronger. An extended synaptic plasticity rule could therefore include a case in which an error complex spike received in the absence of a recent perturbation spike has a neutral plasticity effect. There is currently little data on which to base a detailed implementation, although we note that in our experiments, bursts of 6 climbing fibre stimuli—a strong complex spike—resulted in only a modest LTP compared to the GP_ protocol.

A potentially unsatisfactory aspect of our simulations was the time taken to learn. Of the order of 50,000 iterations were required to optimise the 400 independent variables of the cerebellar output. Stochastic gradient descent is inherently slow, since just one or a few of those variables can be perturbed in one movement realisation, and the weight changes are furthermore individually small. Before considering possible acceleration methods, we note that some motor behaviours are repeated huge numbers of times. An obvious example is locomotion. Thus, public health campaigns in vogue in the USA at the time of writing aim for people to take 10,000 steps per day. So, clearly, a target of a few hundred thousand steps could be achieved in a matter of days or weeks.

Part of the slowness of learning results from the conflicting pressures on the plastic weight changes. Large changes allow rapid learning, but could prevent accurate optimisation. An obvious extension of our implementation that would resolve this conflict would be to allow large plastic changes far from the optimum but to reduce them as the optimum is approached. The information required to do this is available as the net drive (error excitation − cancellation inhibition) to the olivary neurones at the time of emission of an error complex spike. If the drive is strong, one can imagine a long burst of action potentials being emitted. There is in vitro (*Mathy et al., 2009*) and in vivo (*Yang and Lisberger, 2014*; *Rasmussen et al., 2013*) evidence that climbing fibre burst length can influence plasticity in Purkinje cells. It seems possible that the same could be true in the cerebellar nuclei (or alternative plastic site in the subtraction of the average error). However, the above mechanism for adapting learning rates would only work directly in the LTD direction, since olivary cells cannot signal the strength of a net inhibition when no error complex spike is emitted.

A mechanism that could regulate the speed of learning in both LTP and LTD directions would be to target perturbations to the time points where they would be most useful—those with the greatest errors. This might be achieved by increasing the probability (and possibly the strength) of the perturbation complex spikes shortly before strongly unbalanced (excitatory or inhibitory) inputs to the olive. This process offers a possible interpretation for various observations of complex spikes occurring before error evaluation: related to movements (*Bauswein et al., 1983*; *Kitazawa et al., 1998*) or triggered by conditioned stimuli (*Rasmussen et al., 2014*; *Ohmae and Medina, 2015*).

Movement-specific adaptations of the learning rates could provide an explanation for the phenomenon of 'savings', according to which relearning a task after extinction occurs at a faster rate than the initial learning. The adaptations could plausibly be maintained during extinction and therefore persist until the relearning phase. These adaptations could appear to represent memories of previous errors (*Herzfeld et al., 2014*).

Finally, the output convergence solution we proposed above for the bicycle problem could also reflect a parallelisation strategy enabling the computations involved in stochastic gradient descent to be scaled from the small circuit we have simulated to the whole cerebellum. As mentioned above, this might involve one of the above schemes for adjusting learning rates in a way that would help plasticity in regions with 'useful' error information to dominate changes in those without.

Insight into learning in other brain regions

We believe that our proposed implementation of stochastic gradient descent offers possible insight into learning processes in other brain regions.

To date, the most compelling evidence for a stochastic gradient descent mechanism has been provided in the context of the acquisition of birdsong. A specific nucleus, the 'LMAN' has been shown to be responsible for song variability during learning and also to be required for learning (*Doya and Sejnowski, 1988*; *Olveczky et al., 2005*). Its established role is therefore analogous to our perturbation complex spike. Our suggestion that the same input signals both perturbation and error change may also apply in the birdsong context, where it would imply that LMAN also assumes the role of determining the sign of plasticity at the connections it perturbed. However, such an idea has for now not been examined and there is as yet a poor understanding of how the trial song is evaluated and of the mechanism for transmitting that information to the adaptive site; indeed the adaptive site itself has not been identified unequivocally.

We see a stronger potential analogy with our mechanism of stochastic gradient descent in the learning of reward-maximising action sequences by the basal ganglia. Under the stimulus of cortical inputs, ensembles of striatal medial spiny neurones become active, with the resulting activity partition determining the actions selected by disinhibition of central pattern generators through inhibition of the globus pallidus (*Grillner et al., 2005*). It is thought that the system learns to favour the actions that maximise the (discounted) reward, which is signalled by activity bursts in dopaminergic midbrain neurones and phasic release of dopamine, notably in the striatum itself (*Schultz, 1986*). This has been argued (*Schultz et al., 1997*) to reflect *reinforcement learning* or more specifically *temporal difference learning* (*Sutton and Barto, 1998*).

We note that temporal difference learning can be decomposed into two problems: linking actions to potential future rewards and a gradient ascent to maximise reward. In respect of the gradient ascent, we note that dopamine has a second, very well-known action in the striatum: it is necessary for the initiation of voluntary movements, since reduction of dopaminergic input to the striatum is the cause of Parkinson's disease, in which volitional movement is severely impaired. The key point is to combine the two roles of the dopaminergic system in the initiation of movement and in signalling reward. The initiation of movement by dopamine, which could contribute to probabilistic action selection, would be considered analogous to our perturbation complex spike and could create an eligibility trace. A subsequent reward signal would result in plasticity of eligible synapses reinforcing the selection of that action, with possible sites of plasticity including the cortico-striatal synapses that were successful in exciting an ensemble of striatal neurones (Rui Costa made a similar suggestion at the 5th Colloquium of the Institut du Fer à Moulin, Paris, 2014). This would constitute a mechanism of gradient ascent analogous to that we have proposed for gradient descent in the cerebellar system.

Of particular interest is whether correct optimisation also involves a mechanism for subtracting the average error in order to extract gradient information. Such subtraction would be entirely consistent with the reports of midbrain dopaminergic neurones responding more weakly to expected rewards and responding with sub-baseline firing to omission of a predicted reward. These phenomena moreover appear to involve an adaptive inhibitory mechanism (*Eshel et al., 2015*). These observations could be interpreted as a subtraction of the average reward by a process analogous to that we propose for the extraction of the change of error $\delta\mathcal{E}$.

Conclusion

We have proposed a complete and plausible mechanism of stochastic gradient descent in the cerebellar system, in which the climbing fibre perturbs movements, creates an eligibility trace, signals error changes and guides plasticity at the sites of perturbation. We verify predicted plasticity rules that contradict the current consensus and highlight the importance of studying plasticity under physiological conditions. The gradient descent requires extraction of the change of error and we propose an adaptive inhibitory mechanism for doing this via cancellation of the average error. Our implementation of stochastic gradient descent suggests the operation of an analogous mechanism (of gradient ascent) in the basal ganglia initiated and rewarded by dopaminergic signalling.

Materials and methods

Electrophysiology

Animal experimentation methods were authorised by the 'Charles Darwin N°5' ethics committee (authorisation no. 4445). Adult female C57Bl/6 mice (2–5 months old) were anesthetised with isoflurane (Nicholas Piramal Ltd, India) and killed by decapitation. The cerebellum was rapidly dissected into cold solution containing (in mM): 230 sucrose, 26 $NaHCO_3$, 3 KCl, 0.8 $CaCl_2$, 8 $MgCl_2$, 1.25 NaH_2PO_4, 25 D-glucose supplemented with $50\,\mu M$ D-APV to protect the tissue during slicing. $300\,\mu m$ sagittal slices were cut in this solution using a Campden Instruments 7000 smz and stored at $32°C$ in a standard extracellular saline solution containing (in mM): 125 NaCl, 2.5 KCl, 1.5 $CaCl_2$, 1.8 $MgCl_2$, 1.25 NaH_2PO_4, 26 $NaHCO_3$ and 25 D-glucose, bubbled with 95 % O_2 and 5 % CO_2 (pH 7.4). Slices were visualised using an upright microscope with a 40 X, 0.8 NA water-immersion objective and infrared optics (illumination filter 750 ± 50 nm). The recording chamber was continuously perfused at a rate of 4–6 ml min^{-1} with a solution containing (mM): 125 NaCl, 2.5 KCl, 1.5 $CaCl_2$, 1.8 $MgCl_2$, 1.25 NaH_2PO_4, 26 $NaHCO_3$, 25 D-glucose and 10 tricine, a Zn^{2+} buffer (*Paoletti et al., 1997*), bubbled with 95 % O_2 and 5 % CO_2 (pH 7.4). Patch pipettes had resistances in the range 2–4 MΩ with the internal solutions given below. Unless otherwise stated, cells were voltage clamped at −70 mV in the whole-cell configuration. Voltages are reported without correction for the junction potential, which was about 10 mV (so true membrane potentials were more negative than we report). Series resistances were 4–10 MΩ and compensated with settings of ~90 % in a Multiclamp 700B amplifier (Molecular Devices). Whole-cell recordings were filtered at 2 kHz and digitised at 10 kHz. Experiments were performed at 32–34°C. The internal solution contained (in mM): 128 K-gluconate, 10 HEPES, 4 KCl, 2.5 K_2HPO_4, 3.5 Mg-ATP, 0.4 Na_2-GTP, 0.5 L-(–)-malic acid, 0.008 oxaloacetic acid, 0.18 α-ketoglutaric acid, 0.2 pyridoxal 5'-phosphate, 5 L-alanine, 0.15 pyruvic acid, 15 L-glutamine, 4 L-asparagine, 1 reduced L-glutathione, 0.5 NAD^+, 5 phosphocreatine, 1.9 $CaCl_2$, 1.5 $MgCl_2$, 0.1 $K_{3.8}$EGTA. Free $[Ca^{2+}]$ was calculated with Maxchelator (C. Patton, Stanford) to be 120 nM. Chemicals were purchased from Sigma-Aldrich, D-APV from Tocris.

Recordings were made in the vermis of lobules three to eight of the cerebellar cortex. Granule cell EPSCs were elicited with stimulation in the granule cell layer with a glass pipette of tip diameter 8–12 μm filled with HEPES-buffered saline. Climbing fibre electrodes had $\sim 2\,\mu$m diameter and were also positioned in the granule cell layer. Images were taken every 5 min; experiments showing significant slice movement (>20 μm) were discarded. Stimulation intensity was fixed at the beginning of the experiment (1–15 V; 50–200 μs) and maintained unchanged during the experiment.

Analysis

No formal power calculation to determine the sample sizes was performed for the plasticity experiments. Recordings were analysed from 75 cells in 55 animals. Slices were systematically changed after each recording. Animals supplied 1–3 cells to the analysis, usually with rotating induction protocols, but there was no formal randomisation or blocking and there was no blinding. An initial analysis focusing on between-protocol comparisons was performed when $n = 33$ cells were retained for analysis (criteria detailed below). Because the p-value for the GPE vs. G_E difference was close to a $p < 0.001$ threshold (*Colquhoun, 2014*) and the confidence interval for the GPE plasticity ratio was narrower than for the others, acquisition for that protocol was reduced and additional experiments prioritised the G_ _, GP_ and G_E protocols. Experiments were ultimately halted for external reasons, when $n = 58$ cells were retained. The two-stage analysis is a form of multiple comparison; in

consequence, the output of the statistical test on the final data has been corrected by doubling the indicated p-values in *Table 1*.

Inspection of acquired climbing fibre responses revealed some failures of stimuli after the first in a fraction of cells, presumably because the second and subsequent stimuli at short intervals fell within the relative refractory period. As a complex spike was always produced these cells have been included in our analysis, but where individual data are displayed, we identify those cells in which failures of secondary climbing fibre stimuli were observed before the end of the induction period.

Analysis made use of a modular Python framework developed in house. Analysis of EPSC amplitudes for each cell began by averaging all of the EPSCs acquired to give a smooth time course. The time of the peak of this 'global' average response was determined. Subsequent measurement of amplitudes of other averages or of individual responses was performed by averaging the current over 0.5 ms centred on the time of the peak of the global average. The baseline calculated over 5 ms shortly before the stimulus artefact was subtracted to obtain the EPSC amplitude. Similar analyses were performed for both EPSCs of the paired-pulse stimulation.

Individual EPSCs were excluded from further analysis if the baseline current in the sweep exceeded -1 nA at -70 mV. Similarly, the analysis automatically excluded EPSCs in sweeps in which the granule cell stimuli elicited an action potential in the Purkinje cell (possibly through antidromic stimulation of its axon or through capacitive coupling to the electrode). However, during induction, in current clamp, such spikes were accepted. For displaying time series, granule cell responses were averaged in bins of 2 min.

The effects of series resistance changes in the Purkinje cell were estimated by monitoring the transient current flowing in response to a voltage step. The amplitude of the current 2 ms after the beginning of the capacity transient was measured. We shall call this the 'dendritic access'. Modelling of voltage-clamp EPSC recordings in a two-compartment model of a Purkinje cell (*Llano et al., 1991*) suggests that this measure is approximately proportional to EPSC amplitude as the series resistance changes over a reasonable range (not shown). It therefore offers a better estimate of the effect on the EPSC amplitude of series resistance changes than would the value of the series resistance (or conductance), which is far from proportional to EPSC amplitude. Intuitively, this can be seen to arise because the EPSC is filtered by the dendritic compartment and the measure relates to the dendritic component of the capacitive transient, whereas the series resistance relates to the somatic compartment. We therefore calculated R_{res}, the ratio of the dendritic access after induction (when plasticity was assessed) to the value before induction, in order to predict the changes the EPSC amplitude arising from changes of series resistance.

Because we elicited EPSCs using constant-voltage stimulation, variations of the resistance of the tip of the stimulating electrode (for instance if cells are drawn into it) could alter the stimulating current flow in the tissue. We monitored this by measuring the amplitude of the stimulus artefact. Specifically, we calculated R_{stim}, the after/before ratio of the stimulation artefact amplitude.

We then used a robust linear model to examine the extent to which changes of series resistance or apparent stimulation strength could confound our measurements of plasticity, which we represented as the after/before ratio of EPSC amplitudes R_{EPSC}; the model (in R syntax) was:

$$R_{EPSC} \sim \text{protocol} + R_{res} + R_{stim}$$

This showed that series resistance changes, represented by R_{res}, had a significant influence (t-value 2.30, 69 degrees of freedom) with a slope close to the predicted unity (1.14). In contrast, changes of the stimulus artefact had no predictive value (slope -0.003, t-value -0.01).

Extending our modelling using a mixed-effect model (call 'lmer' in the lme4 R package) to include animals as a random effect did not indicate that it was necessary to take into account any between-animal variance, as this was reported to be zero. On this basis, we consider each cell recorded to be the biological replicate.

We did not wish to rely on the parametric significance tests of the linear model for comparing the plasticity protocols (although all of comparisons we report below as significant were also significant in the model). Instead, we equalised the dendritic filtering and stimulation changes between groups by eliminating those cells in which R_{res} or R_{stim} differed by more than 20% from the mean values for all cells (0.94 $\pm$ 0.10 and 1.01 $\pm$ 0.19, respectively; mean $\pm$ sd, $n = 75$). After this operation, which eliminated 17 cells out of 75 leaving 58 (from 47 animals), the mean ratios varied by only a few

percent between groups (ranges 5 % and 2 % for R_{res} and R_{stim}, respectively) and would be expected to have only a minimal residual influence. Normalising the R_{EPSC}s of the trimmed groups by R_{res} did not alter the conclusions presented below. Note that after this trimming, the remaining changes of R_{res} imply that all EPSC amplitudes after induction were underestimated by about 6 % relative to those at the beginning of the recording. The differences of R_{EPSC} between induction protocols were evaluated statistically using two-tailed nonparametric tests implemented by the 'wilcox.test' command in R (*R Core Team, 2013*).

95 % confidence limits were calculated using R bootstrap functions ('BCa' method). For confidence limits of differences between means, stratified resampling was used.

A small additional series of experiments of 16 cells from 16 animals of which 12 were retained by similar criteria to those explained above (except that series resistance changes had to be evaluated using the current amplitude 1 ms into the response to the voltage step) were analysed separately and underlie the plasticity data in *Figure 6—figure supplement 1*.

Simulation methods

Network simulations were designed as follows (see also diagram in *Figure 12*). A total of $S \times L$ Purkinje cells were placed on a rectangular grid of extent S in the sagittal plane and width L in the lateral direction. The activity of each Purkinje cell during a 'movement' was characterised by its firing rate in T time bins $\{t = 1, \cdots, T\}$.

$$\{PC_{s,l}(t), s = 1, \cdots, S, l = 1, \cdots, L\} \tag{1}$$

Purkinje cells contacted projection neurones (*PN*) and nucleo-olivary neurones (*NO*) in the cerebellar nuclei that contained L cells of each type. The activities of both types of nuclear cell were also characterised by their firing rates in the different time bins, $\{PN_l(t), l = 1, \cdots, L\}$ and $\{NO_l(t), l = 1, \cdots, L\}$. Mossy fibres, granule cells (parallel fibres) and molecular layer interneurons were subsumed into a single cell type M (for mossy fibre) with N cells restricted to each row of L Purkinje cells, with a total of $N \times S$ mossy fibres. Mossy fibre activity was represented in a binary manner, $M_{i,s}(t) = 0$ or 1, with $\{i = 1, \cdots, N\}$ and $\{s = 1, \cdots, S\}$.

The connectivity was chosen such that all Purkinje cells in the sagittal 'column' l projected to the l-th nuclear projection neurone with identical inhibitory (negative) weights, as well as to the l-th nucleo-olivary neurone with potentially different, but also identical inhibitory weights. Mossy fibres were chosen to project to Purkinje cells with groups of N fibres at a given sagittal position s, contacting cells on the row of L Purkinje cells at the sagittal position with probability 1/2. The activity of Purkinje and cerebellar nuclear cells in the absence of climbing fibre activity was thus

$$PC_{s,l}(t) = \Phi\left(\sum_{i=1,N} w_{i,s}^l \sigma_{i,s}^l M_{i,s}(t)\right) \tag{2}$$

$$PN_l(t) = \Phi\left(\sum_{i=1,N;s=1,S} u_{i,s}^l M_{i,s}(t) + u_{(PC \to PN)} \sum_{s=1,S} PC_{s,l}(t)\right) \tag{3}$$

$$NO_l(t) = \Phi\left(\sum_{i=1,N;s=1,S} v_{i,s}^l M_{i,s}(t) + u_{(PC \to NO)} \sum_{s=1,S} PC_{s,l}(t)\right), \tag{4}$$

where $\sigma_{i,s}^l$ enforces the 1/2 probability of connection between a Purkinje cell and a parallel fibre that traverses its dendritic tree: $\sigma_{i,s}^l$ is equal to 1 or 0 with probability 1/2, independently drawn for each triple (i,s,l). The f-I curve Φ was taken to be a saturating threshold linear function, $\{\Phi(x) = 0$ for $x < 0$, $\Phi(x) = x$ for $0 < x < r_{max}$ and $\Phi(x) = r_{max}$ for $x > r_{max}\}$. The weights $u_{i,s}^l$ were non-plastic and chosen such that a given mossy fibre (i,s) contacted a single projection neurone and a single nucleo-olivary neurone among the L possible ones of each type. In other words, for each mossy fibre index (i,s), a number $l_{(i,s)}$ was chosen at random with uniform probability among the L numbers $\{1, \cdots, L\}$ and the weights $\left\{u_{i,s}^l, l = 1, \cdots, l\right\}$ were determined as
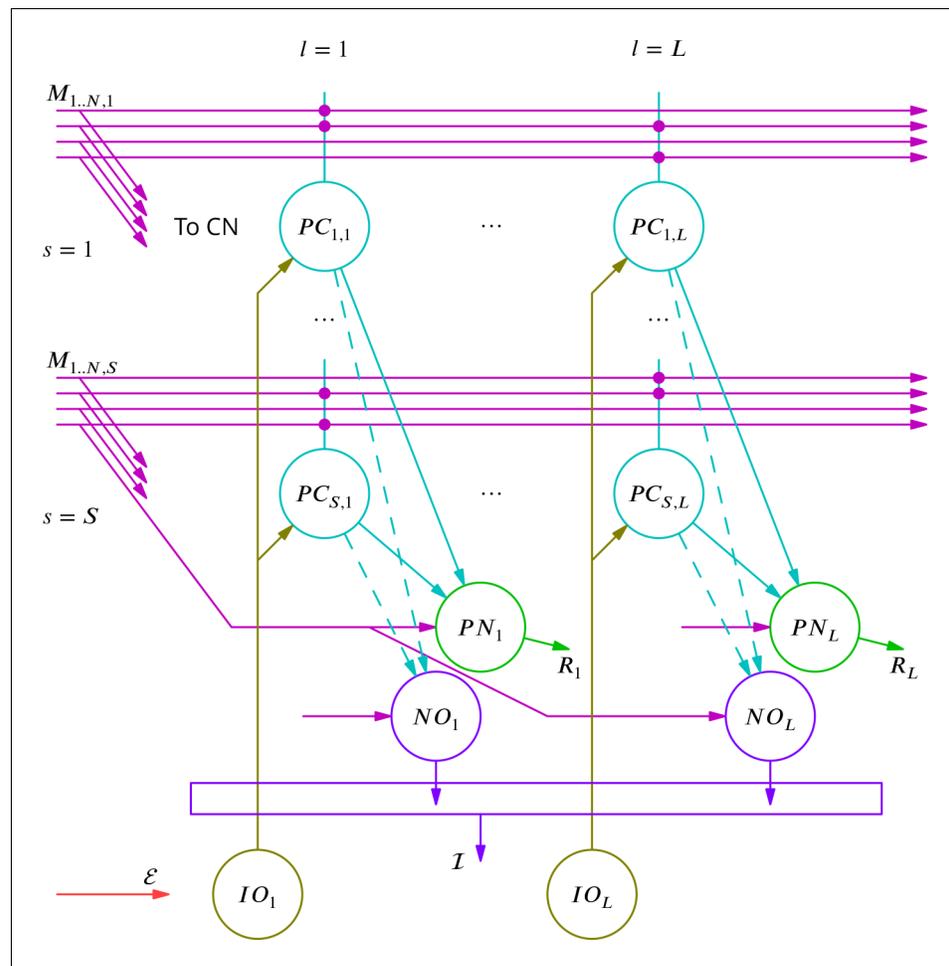
Figure 12. Diagram of the simulated network model. The details are explained in the text.
DOI: https://doi.org/10.7554/eLife.31599.015

$$u^l_{i,s} = \delta_{l,l_{(i,s)}} u_{(M \to PN)}, \tag{5}$$

where $u_{(M \to PN)}$ was a constant. The weights $v^l_{i,s}$ and $w^l_{i,s}$ were plastic and followed the learning dynamics described below (see *Equations 11 and 13*).

The learning task itself consisted of producing, in response to $\mu = 1, \cdots, p$ spatiotemporal patterns of mossy fibre inputs $M^\mu_{i,s}(t)$, the corresponding output target rates of the projection neurones $\{R^\mu_l(t) \, l = 1, \cdots, L, \, t = 1, ..., T\}$. For each pattern μ, the inputs were obtained by choosing at random with uniform probability $NS/2$ active fibres. For each active fibre (i, s), a time bin $t^\mu_{(i,s)}$ was chosen at random with uniform probability and the activity of the fibre was set to one in this time bin, $M^\mu_{(i,s)}(t) = \delta_{t,t^\mu_{(i,s)}}$. The activity was set to zero in all time bins for the $NS/2$ inactive fibres. The target rates where independently chosen with uniform probability between 0 and $2\bar{r}_D$ for each projection neurone in each pattern μ, where $\bar{r}_D$ is the desired average firing rate for both projection and nucleo-olivary neurones in the cerebellar nuclei.

The olivary neurones were not explicitly represented. It was assumed that the $L \times S$ Purkinje cells were contacted by L climbing fibres with one climbing fibre contacting the S Purkinje cells at a given lateral position.

The learning algorithm then proceeded as follows. Patterns $\{M^\mu_{i,s}(t)\}$ were presented sequentially. After pattern μ was chosen, perturbations of Purkinje cell firing by complex spikes were

generated as follows. The probability that each climbing fibre emitted a perturbation complex spike was taken to be ρ per pattern presentation; when a climbing fibre was active, it was considered to perturb the firing of its Purkinje cells in a single time bin chosen at random. Denoting by $\eta_l(t) = 1$ that climbing fibre l had emitted a spike in time bin t (and $\eta_l(t) = 0$ when there was no spike), the S firing rates of the Purkinje cell at position l (see *Equation 3*) were taken to be

$$PC_{s,l}(t) = \Phi\left(\sum_{i=1,N} w_{i,s}^l \sigma_{i,s}^l M_{i,s}(t) + \eta_l(t)A\right), s = 1, \cdots, S, \tag{6}$$

where A defines the amplitude of the complex-spike perturbation of Purkinje cell firing.

Given the pattern μ and the firing of Purkinje cells (*Equation 6*), the activities of cerebellar nuclear neurones were given by *Equations 3 and 4*. The current 'error' for pattern/movement μ was quantified by the average distance of the projection neurones' activity from their target rates

$$\mathcal{E}^\mu = \frac{1}{LT} \sum_{l=1,L; t=1,T} |PN_l^\mu(t) - R_l^\mu(t)|. \tag{7}$$

The learning step after the presentation of pattern μ was determined by the comparison (not explicitly implemented) in the olivary neurones between the excitation $\mathcal{E}^\mu$ and the inhibition $\mathcal{I}^\mu$ coming from the discharges of nucleo-olivary neurones with

$$\mathcal{I}^\mu = \frac{1}{LT} \sum_{l=1,L; t=1,T} NO_l^\mu(t). \tag{8}$$

An error complex spike was propagated to all Purkinje cells after a 'movement' when the olivary activity $IO = \mathcal{E}^\mu - \mathcal{I}^\mu$ was positive. Accordingly, modifications of the weights of mossy fibre synapses on perturbed Purkinje cells (w) and on nucleo-olivary neurones (v) were determined after presentation of pattern μ by the sign c of IO, $c = \text{sign}(IO)$, as,

$$w_{i,s}^l \rightarrow w_{i,s}^l - \alpha_w c \sum_{t=1,T} \eta_l(t) M_{i,s}^\mu(t) \tag{9}$$

$$v_{i,s}^l \rightarrow \left[v_{i,s}^l + \alpha_v c \sum_{t=1,T} M_{i,s}^\mu(t)\right]_+ \tag{10}$$

where the brackets served to enforce a positivity constraint on the weights $v_{i,s}^l$ ($[x]_+ = x, x > 0$ and $[x]_+ = 0, x < 0$).

In the reported simulations, the non-plastic weights in *Equations 3–5* were identical and constant for all synapses of a given type:

$$u_{(M \rightarrow PN)} = 4TL\bar{r}_D / NS \tag{11}$$

$$u_{(PC \rightarrow PN)} = -\bar{r}_D / (\bar{r}_{PC} S) \tag{12}$$

$$u_{(PC \rightarrow NO)} = q u_{(PC \rightarrow PN)} = -q\bar{r}_D / (\bar{r}_{PC} S), \tag{13}$$

where q describes the strength of the $PC \rightarrow NO$ connection relative to the $PC \rightarrow PN$ connection. The initial weights of the plastic synapses were drawn from uniform distributions such that initial firing rates were on average given by $\overline{PC}_0 = \bar{r}_{PC}$ and $\overline{NO}_0 = 15$ Hz for Purkinje cells and nucleo-olivary neurones, respectively. The latter value was chosen such that the expected initial mismatch between inhibition $\mathcal{I}$ and excitation $\mathcal{E}$ for random pairings of initial PN firing rates and target rates R was of the order of 1 Hz. Consequently, for $M \rightarrow PC$ synapses $w_{i,s}^l$ was drawn from $[0, 8T\bar{r}_{PC}/N]$ and for $M \rightarrow NO$ synapses, weights $v_{i,s}^l$ were drawn from $\left[0, (\overline{NO}_0 + q\bar{r}_D)4TL/(NS)\right]$. These weights furthermore ensured initial average firing rates close to $\bar{r}_D$ in the nuclear projection neurones. Target rates R are uniformly distributed between 0 Hz and $2\bar{r}_D$.

The parameters used in the reported simulation (full and simplified) are provided in *Table 2*.

In the following we describe the simplified version of the model without perturbations used to highlight shortcomings of the current Marr-Albus-Ito theory. M, PC, and PN cell types are defined as above with identical synaptic connections and identical initial weights. The learning rules are different as no comparison between an estimated error (inhibition) and the global error (excitation) is made. We therefore do not consider the population of NO cells in this version. Two methods of calculating the error were tested. In the first, the error conveyed by the climbing fibres contains information about the sign of the movement error and is given by

$$\mathcal{E}^{\mu} = \frac{1}{LT} \sum_{l=1,L;t=1,T} \left(R_l^{\mu}(t) - PN_l^{\mu}(t) \right). \tag{14}$$

In the second, the absolute values according to **Equation 7** were used.

The synaptic weight changes are as follows. Whenever the error is positive, the synaptic weights of active $M \to PC$ connections are decreased by an amount α_w, while otherwise there is a slow positive drift of $M \to PC$ synaptic weights with increments β_w:

$$w_{i,s}^l \to \begin{cases} w_{i,s}^l - \alpha_w \sum_{t=1,T} M_{i,s}^{\mu}(t) & \text{if } \mathcal{E}^{\mu} > 0 \\ w_{i,s}^l + \beta_w & \text{else.} \end{cases} \tag{15}$$

In order to show relaxation of the average of PN firing rates to the average of the target rates R, we chose different initial distributions of uniformly distributed target rates with averages between 10 Hz and 50 Hz.

We show in the main text that the four key parameters governing convergence of our learning algorithm are the learning rates of mossy fibre–Purkinje cell synapses α_w and mossy fibre–nucleo-olivary neurone synapses α_v, as well as the probability of a perturbation complex spike occurring in a given movement in a given cell ρ, and the resulting amplitude of the perturbation of Purkinje cell firing A. Varying each of these $\pm$ 10 % individually altered the final error (1.4 Hz, averaged over trials 55,000–60,000) by at most 7%, indicating that this final output was not ill-conditioned or finely tuned with respect to these parameters.

The simulation was coded in Python.

Table 2. Parameters used in the simulation shown in **Figure 12**.

Parameter	Symbol	Value
Sagittal extent	S	10
Lateral extent	L	40
Time bins per movement	T	10
Mossy fibres per sagittal position	N	2000
Maximum firing rate (all neurones)	r_{max}	300 Hz
pCS probability per cell per movement	ρ	0.03
Amplitude of Purkinje cell firing perturbation	A	2 Hz
Learning rate of $M \to PC$ synapses	α_w	0.02
Learning rate of $M \to NO$ synapses	α_v	0.0002
Mean Purkinje cell firing rate	$\bar{r}_{PC}$	50 Hz
Mean firing rate for nuclear neurones	$\bar{r}_D$	30 Hz
$M \to PN$ synaptic weight	$u_{(M \to PN)}$	2.4
$PC \to PN$ synaptic weight	$u_{(PC \to PN)}$	−0.06
$PC \to NO$ relative synaptic weight	q	0.5
Weight increment of $M \to PC$ drift (simple model)	β_w	0.002

Acknowledgements

Additional information

Author contributions

Guy Bouvier, Data curation, Investigation, Methodology, Writing—original draft, Writing—review and editing, Performed all experiments shown; Johnatan Aljadeff, Jonas Ranft, Software, Formal analysis, Writing—original draft, Writing—review and editing; Claudia Clopath, Writing—review and editing, Contributed to Initial theoretical exploration; Célian Bimbard, Investigation, Performed pilot experiments and suggested implementation of tracking plasticity; Antonin Blot, Software; Jean-Pierre Nadal, Formal analysis; Nicolas Brunel, Vincent Hakim, Conceptualization, Software, Formal analysis, Supervision, Writing—original draft, Writing—review and editing; Boris Barbour, Conceptualization, Data curation, Software, Formal analysis, Supervision, Funding acquisition, Methodology, Writing—original draft, Project administration, Writing—review and editing, Performed all data analysis shown

Author ORCIDs

Guy Bouvier [iD] https://orcid.org/0000-0002-6160-7186
Johnatan Aljadeff [iD] http://orcid.org/0000-0002-7145-0514
Claudia Clopath [iD] http://orcid.org/0000-0003-4507-8648
Célian Bimbard [iD] http://orcid.org/0000-0002-6380-5856
Jonas Ranft [iD] http://orcid.org/0000-0002-7843-7443
Antonin Blot [iD] https://orcid.org/0000-0002-1546-3927
Jean-Pierre Nadal [iD] http://orcid.org/0000-0003-0022-0647
Nicolas Brunel [iD] https://orcid.org/0000-0002-2272-3248
Vincent Hakim [iD] http://orcid.org/0000-0001-7505-8192
Boris Barbour [iD] http://orcid.org/0000-0002-1911-0539

Ethics

Animal experimentation: Animal experimentation methods were performed according to authorisation 04445.02 granted by the 'Charles Darwin N°5' ethics committee.

Decision letter and Author response

Decision letter https://doi.org/10.7554/eLife.31599.025
Author response https://doi.org/10.7554/eLife.31599.026

Additional files

Supplementary files

• Transparent reporting form
DOI: https://doi.org/10.7554/eLife.31599.017

Data availability

Source data, analysis/simulation scripts and software libraries have been depositied at the Zenodo repository.

The following datasets were generated:

Author(s)	Year	Dataset title	Dataset URL	Database and Identifier
Guy Bouvier, Johnatan Aljadeff, Claudia Clopath, Célian Bimbard, Jonas Ranft, Antonin Blot, Jean-Pierre Nadal, Nicolas Brunel, Vincent Hakim, Boris Barbour	2018	Cerebellar learning using perturbations: data, analysis/simulation scripts	https://doi.org/10.5281/zenodo.1481929	Zenodo, 10.5281/zenodo.1481929
Guy Bouvier, Johnatan Aljadeff, Claudia Clopath, Célian Bimbard, Jonas Ranft, Antonin Blot, Jean-Pierre Nadal, Nicolas Brunel, Vincent Hakim, Boris Barbour	2018	Cerebellar learning using perturbations: software libraries	https://doi.org/10.5281/zenodo.1481925	Zenodo, 10.5281/zenodo.1481925

References

Albus JS. 1971. A theory of cerebellar function. *Mathematical Biosciences* **10**:25–61. DOI: https://doi.org/10.1016/0025-5564(71)90051-4

Apps R, Hawkes R. 2009. Cerebellar cortical organization: a one-map hypothesis. *Nature Reviews Neuroscience* **10**:670–681. DOI: https://doi.org/10.1038/nrn2698, PMID: 19693030

Arima-Yoshida F, Watabe AM, Manabe T. 2011. The mechanisms of the strong inhibitory modulation of long-term potentiation in the rat dentate gyrus. *European Journal of Neuroscience* **33**:1637–1646. DOI: https://doi.org/10.1111/j.1460-9568.2011.07657.x, PMID: 21535245

Barmack NH, Hess DT. 1980. Eye movements evoked by microstimulation of dorsal cap of inferior olive in the rabbit. *Journal of Neurophysiology* **43**:165–181. DOI: https://doi.org/10.1152/jn.1980.43.1.165

Barto AG, Sutton RS, Anderson CW. 1983. Neuronlike adaptive elements that can solve difficult learning control problems. IEEE Transactions on Systems, Man, and Cybernetics 834–846. DOI: https://doi.org/10.1109/TSMC.1983.6313077

Bauswein E, Kolb FP, Leimbeck B, Rubia FJ. 1983. Simple and complex spike activity of cerebellar purkinje cells during active and passive movements in the awake monkey. *The Journal of Physiology* **339**:379–394. DOI: https://doi.org/10.1113/jphysiol.1983.sp014722, PMID: 6887028

Bazzigaluppi P, De Gruijl JR, van der Giessen RS, Khosrovani S, De Zeeuw CI, de Jeu MT. 2012a. Olivary subthreshold oscillations and burst activity revisited. *Frontiers in Neural Circuits* **6**:91. DOI: https://doi.org/10.3389/fncir.2012.00091, PMID: 23189043

Bazzigaluppi P, Ruigrok T, Saisan P, De Zeeuw CI, de Jeu M. 2012b. Properties of the nucleo-olivary pathway: an in vivo whole-cell patch clamp study. *PLOS ONE* **7**:e46360. DOI: https://doi.org/10.1371/journal.pone.0046360, PMID: 23029495

Bengtsson F, Ekerot CF, Jörntell H. 2011. In vivo analysis of inhibitory synaptic inputs and rebounds in deep cerebellar nuclear neurons. *PLOS ONE* **6**:e18822. DOI: https://doi.org/10.1371/journal.pone.0018822, PMID: 21552556

Best AR, Regehr WG. 2009. Inhibitory regulation of electrically coupled neurons in the inferior olive is mediated by asynchronous release of GABA. *Neuron* **62**:555–565. DOI: https://doi.org/10.1016/j.neuron.2009.04.018, PMID: 19477156

Blazquez PM, Hirata Y, Highstein SM. 2004. The vestibulo-ocular reflex as a model system for motor learning: what is the role of the cerebellum? *The Cerebellum* **3**:188–192. DOI: https://doi.org/10.1080/14734220410018120, PMID: 15543809

Boele HJ, Koekkoek SK, De Zeeuw CI, Ruigrok TJ. 2013. Axonal sprouting and formation of terminals in the adult cerebellum during associative motor learning. *Journal of Neuroscience* **33**:17897–17907. DOI: https://doi.org/10.1523/JNEUROSCI.0511-13.2013, PMID: 24198378

Bouvier G, Higgins D, Spolidoro M, Carrel D, Mathieu B, Léna C, Dieudonné S, Barbour B, Brunel N, Casado M. 2016. Burst-Dependent bidirectional plasticity in the cerebellum is driven by presynaptic NMDA receptors. *Cell Reports* **15**:104–116. DOI: https://doi.org/10.1016/j.celrep.2016.03.004, PMID: 27052175

Brunel N, Hakim V, Isope P, Nadal JP, Barbour B. 2004. Optimal information storage and the distribution of synaptic weights: perceptron versus purkinje cell. *Neuron* **43**:745–757. DOI: https://doi.org/10.1016/j.neuron.2004.08.023, PMID: 15339654

Caddy KWT, Biscoe TJ. 1976. The number of Purkinje cells and olive neurones in the normal and Lurcher mutant mouse. *Brain Research* **111**:396–398. DOI: https://doi.org/10.1016/0006-8993(76)90783-6

Campbell NC, Hesslow G. 1986. The secondary spikes of climbing fibre responses recorded from purkinje cell axons in cat cerebellum. *The Journal of Physiology* **377**:225–235. DOI: https://doi.org/10.1113/jphysiol.1986.sp016183, PMID: 3795088

Catz N, Dicke PW, Thier P. 2005. Cerebellar complex spike firing is suitable to induce as well as to stabilize motor learning. *Current Biology* **15**:2179–2189. DOI: https://doi.org/10.1016/j.cub.2005.11.037, PMID: 16360681

Chadderton P, Margrie TW, Häusser M. 2004. Integration of quanta in cerebellar granule cells during sensory processing. *Nature* **428**:856–860. DOI: https://doi.org/10.1038/nature02442, PMID: 15103377

Clopath C, Brunel N. 2013. Optimal properties of analog perceptrons with excitatory weights. *PLOS Computational Biology* **9**:e1002919. DOI: https://doi.org/10.1371/journal.pcbi.1002919, PMID: 23436991

Coesmans M, Weber JT, De Zeeuw CI, Hansel C. 2004. Bidirectional parallel fiber plasticity in the cerebellum under climbing fiber control. *Neuron* **44**:691–700. DOI: https://doi.org/10.1016/j.neuron.2004.10.031, PMID: 15541316

Colquhoun D. 2014. An investigation of the false discovery rate and the misinterpretation of p-values. *Royal Society Open Science* **1**:140216. DOI: https://doi.org/10.1098/rsos.140216, PMID: 26064558

Crepel F, Jaillard D. 1991. Pairing of pre- and postsynaptic activities in cerebellar purkinje cells induces long-term changes in synaptic efficacy in vitro. *The Journal of Physiology* **432**:123–141. DOI: https://doi.org/10.1113/jphysiol.1991.sp018380, PMID: 1886056

Dash S, Thier P. 2014. Cerebellum-dependent motor learning: lessons from adaptation of eye movements in primates. *Progress in Brain Research* **210**:121–155. DOI: https://doi.org/10.1016/B978-0-444-63356-9.00006-6, PMID: 24916292

Dayan P, Abbott LF. 2001. *Theoretical Neuroscience*. **806** Cambridge, MA: MIT Press.

De Zeeuw CI, Van Alphen AM, Hawkins RK, Ruigrok TJ. 1997. Climbing fibre collaterals contact neurons in the cerebellar nuclei that provide a GABAergic feedback to the inferior olive. *Neuroscience* **80**:981–986. PMID: 9284054

De Zeeuw CI, Simpson JI, Hoogenraad CC, Galjart N, Koekkoek SK, Ruigrok TJ. 1998. Microcircuitry and function of the inferior olive. *Trends in Neurosciences* **21**:391–400. DOI: https://doi.org/10.1016/S0166-2236(98)01310-1, PMID: 9735947

Dean P, Porrill J, Stone JV. 2002. Decorrelation control by the cerebellum achieves oculomotor plant compensation in simulated vestibulo-ocular reflex. *Proceedings of the Royal Society B: Biological Sciences* **269**:1895–1904. DOI: https://doi.org/10.1098/rspb.2002.2103

Dean P, Porrill J. 2014. Decorrelation learning in the cerebellum: computational analysis and experimental questions. *Progress in Brain Research* **210**:157–192. DOI: https://doi.org/10.1016/B978-0-444-63356-9.00007-8, PMID: 24916293

Devor A, Yarom Y. 2002. Electrotonic coupling in the inferior olivary nucleus revealed by simultaneous double patch recordings. *Journal of Neurophysiology* **87**:3048–3058. DOI: https://doi.org/10.1152/jn.2002.87.6.3048, PMID: 12037207

Doya K, Sejnowski TJ. 1988. A computational model of birdsong learning by auditory experience and auditory feedback. In: Brugge J, Poon P (Eds). *Central Auditory Processing and Neural Modeling*. New York: Plenum Press. p. 77–88.

Eccles JC, Llinás R, Sasaki K. 1966. The excitatory synaptic action of climbing fibres on the purkinje cells of the cerebellum. *The Journal of Physiology* **182**:268–296. DOI: https://doi.org/10.1113/jphysiol.1966.sp007824, PMID: 5944665

Eccles JC, Ito M, Szentagothai J. 1967. *The Cerebellum as a Neuronal Machine*. Berlin Heidelberg: Springer .

Ekerot CF, Jörntell H, Garwicz M. 1995. Functional relation between corticonuclear input and movements evoked on microstimulation in cerebellar nucleus interpositus anterior in the cat. *Experimental Brain Research* **106**:365–376. DOI: https://doi.org/10.1007/BF00231060, PMID: 8983981

Eshel N, Bukwich M, Rao V, Hemmelder V, Tian J, Uchida N. 2015. Arithmetic and local circuitry underlying dopamine prediction errors. *Nature* **525**:243–246. DOI: https://doi.org/10.1038/nature14855, PMID: 26322583

Fiete IR, Fee MS, Seung HS. 2007. Model of birdsong learning based on gradient estimation by dynamic perturbation of neural conductances. *Journal of Neurophysiology* **98**:2038–2057. DOI: https://doi.org/10.1152/jn.01311.2006, PMID: 17652414

Garwicz M, Jorntell H, Ekerot CF. 1998. Cutaneous receptive fields and topography of mossy fibres and climbing fibres projecting to cat cerebellar C3 zone. *The Journal of Physiology* **512**:277–293. DOI: https://doi.org/10.1111/j.1469-7793.1998.277bf.x, PMID: 9729638

Garwicz M, Levinsson A, Schouenborg J. 2002. Common principles of sensory encoding in spinal reflex modules and cerebellar climbing fibres. *The Journal of Physiology* **540**:1061–1069. DOI: https://doi.org/10.1113/jphysiol.2001.013507, PMID: 11986390

Graf W, Simpson JI, Leonard CS. 1988. Spatial organization of visual messages of the rabbit's cerebellar flocculus. II. Complex and simple spike responses of Purkinje cells. *Journal of Neurophysiology* **60**:2091–2121. DOI: https://doi.org/10.1152/jn.1988.60.6.2091, PMID: 3236063

Graupner M, Brunel N. 2007. STDP in a bistable synapse model based on CaMKII and associated signaling pathways. *PLOS Computational Biology* **3**:e221. DOI: https://doi.org/10.1371/journal.pcbi.0030221, PMID: 18052535

Grillner S, Hellgren J, Ménard A, Saitoh K, Wikström MA. 2005. Mechanisms for selection of basic motor programs–roles for the striatum and pallidum. *Trends in Neurosciences* **28**:364–370. DOI: https://doi.org/10.1016/j.tins.2005.05.004, PMID: 15935487

Gutierrez-Castellanos N, Da Silva-Matos CM, Zhou K, Canto CB, Renner MC, Koene LMC, Ozyildirim O, Sprengel R, Kessels HW, De Zeeuw CI. 2017. Motor learning requires purkinje cell synaptic potentiation through activation of AMPA-Receptor subunit GluA3. *Neuron* **93**:409–424. DOI: https://doi.org/10.1016/j.neuron.2016.11.046, PMID: 28103481

Harris CM. 1998. On the optimal control of behaviour: a stochastic perspective. *Journal of Neuroscience Methods* **83**:73–88. DOI: https://doi.org/10.1016/S0165-0270(98)00063-6, PMID: 9765052

Herzfeld DJ, Vaswani PA, Marko MK, Shadmehr R. 2014. A memory of errors in sensorimotor learning. *Science* **345**:1349–1353. DOI: https://doi.org/10.1126/science.1253138, PMID: 25123484

Ito M, Simpson JI. 1971. Discharges in purkinje cell axons during climbing fiber activation. *Brain Research* **31**:215–219. DOI: https://doi.org/10.1016/0006-8993(71)90648-2, PMID: 4328277

Ito M. 1972. Neural design of the cerebellar motor control system. *Brain Research* **40**:81–84. DOI: https://doi.org/10.1016/0006-8993(72)90110-2, PMID: 4338265

Ito M, Shiida T, Yagi N, Yamamoto M. 1974. Visual influence on rabbit horizontal vestibulo-ocular reflex presumably effected via the cerebellar flocculus. *Brain Research* **65**:170–174. DOI: https://doi.org/10.1016/0006-8993(74)90344-8, PMID: 4810171

Ito M, Kano M. 1982. Long-lasting depression of parallel fiber-Purkinje cell transmission induced by conjunctive stimulation of parallel fibers and climbing fibers in the cerebellar cortex. *Neuroscience Letters* **33**:253–258. DOI: https://doi.org/10.1016/0304-3940(82)90380-9, PMID: 6298664

Ito M, Sakurai M, Tongroach P. 1982. Climbing fibre induced depression of both mossy fibre responsiveness and glutamate sensitivity of cerebellar purkinje cells. *The Journal of Physiology* **324**:113–134. DOI: https://doi.org/10.1113/jphysiol.1982.sp014103, PMID: 7097592

Ito M. 1984. *The Cerebellum and Neural Control*. Raven Press.

Jones HC, Keep RF. 1988. Brain fluid calcium concentration and response to acute hypercalcaemia during development in the rat. *The Journal of Physiology* **402**:579–593. DOI: https://doi.org/10.1113/jphysiol.1988.sp017223, PMID: 3236250

Jörntell H, Garwicz M, Ekerot CF. 1996. Relation between cutaneous receptive fields and muscle afferent input to climbing fibres projecting to the cerebellar C3 zone in the cat. *European Journal of Neuroscience* **8**:1769–1779. DOI: https://doi.org/10.1111/j.1460-9568.1996.tb01320.x, PMID: 8921267

Jörntell H, Ekerot CF. 2002. Reciprocal bidirectional plasticity of parallel fiber receptive fields in cerebellar purkinje cells and their afferent interneurons. *Neuron* **34**:797–806. DOI: https://doi.org/10.1016/S0896-6273(02)00713-4, PMID: 12062025

Jörntell H, Ekerot CF. 2003. Receptive field plasticity profoundly alters the cutaneous parallel fiber synaptic input to cerebellar interneurons in vivo. *The Journal of Neuroscience* **23**:9620–9631. DOI: https://doi.org/10.1523/JNEUROSCI.23-29-09620.2003, PMID: 14573542

Jörntell H, Ekerot CF. 2006. Properties of somatosensory synaptic integration in cerebellar granule cells in vivo. *Journal of Neuroscience* **26**:11786–11797. DOI: https://doi.org/10.1523/JNEUROSCI.2939-06.2006, PMID: 17093099

Jörntell H, Ekerot CF. 2011. Receptive field remodeling induced by skin stimulation in cerebellar neurons in vivo. *Frontiers in Neural Circuits* **5**:3. DOI: https://doi.org/10.3389/fncir.2011.00003, PMID: 21427779

Ke MC, Guo CC, Raymond JL. 2009. Elimination of climbing fiber instructive signals during motor learning. *Nature Neuroscience* **12**:1171–1179. DOI: https://doi.org/10.1038/nn.2366, PMID: 19684593

Khaliq ZM, Raman IM. 2005. Axonal propagation of simple and complex spikes in cerebellar purkinje neurons. *Journal of Neuroscience* **25**:454–463. DOI: https://doi.org/10.1523/JNEUROSCI.3045-04.2005, PMID: 15647489

Kim JJ, Krupa DJ, Thompson RF. 1998. Inhibitory cerebello-olivary projections and blocking effect in classical conditioning. *Science* **279**:570–573. DOI: https://doi.org/10.1126/science.279.5350.570, PMID: 9438852

Kimpo RR, Rinaldi JM, Kim CK, Payne HL, Raymond JL. 2014. Gating of neural error signals during motor learning. *eLife* **3**:e02076. DOI: https://doi.org/10.7554/eLife.02076, PMID: 24755290

Kitazawa S, Kimura T, Yin PB. 1998. Cerebellar complex spikes encode both destinations and errors in arm movements. *Nature* **392**:494–497. DOI: https://doi.org/10.1038/33141, PMID: 9548253

Kleim JA, Freeman JH, Bruneau R, Nolan BC, Cooper NR, Zook A, Walters D. 2002. Synapse formation is associated with memory storage in the cerebellum. *PNAS* **99**:13228–13231. DOI: https://doi.org/10.1073/pnas. 202483399, PMID: 12235373

Konishi M. 1965. The role of auditory feedback in the control of vocalization in the white-crowned sparrow. *Zeitschrift Für Tierpsychologie* **22**:770–783.

Lee KH, Mathews PJ, Reeves AM, Choe KY, Jami SA, Serrano RE, Otis TS. 2015. Circuit mechanisms underlying motor memory formation in the cerebellum. *Neuron* **86**:529–540. DOI: https://doi.org/10.1016/j.neuron.2015. 03.010, PMID: 25843404

Lev-Ram V, Makings LR, Keitz PF, Kao JP, Tsien RY. 1995. Long-term depression in cerebellar purkinje neurons results from coincidence of nitric oxide and depolarization-induced Ca2+ transients. *Neuron* **15**:407–415. DOI: https://doi.org/10.1016/0896-6273(95)90044-6, PMID: 7646893

Lev-Ram V, Wong ST, Storm DR, Tsien RY. 2002. A new form of cerebellar long-term potentiation is postsynaptic and depends on nitric oxide but not cAMP. *PNAS* **99**:8389–8393. DOI: https://doi.org/10.1073/pnas. 122206399, PMID: 12048250

Lev-Ram V, Mehta SB, Kleinfeld D, Tsien RY. 2003. Reversing cerebellar long-term depression. *PNAS* **100**: 15989–15993. DOI: https://doi.org/10.1073/pnas.2636935100

Llano I, Marty A, Armstrong CM, Konnerth A. 1991. Synaptic- and agonist-induced excitatory currents of purkinje cells in rat cerebellar slices. *The Journal of Physiology* **434**:183–213. DOI: https://doi.org/10.1113/jphysiol. 1991.sp018465, PMID: 1673717

Llinás R, Mühlethaler M. 1988. Electrophysiology of guinea-pig cerebellar nuclear cells in the in vitro brain stem-cerebellar preparation. *The Journal of Physiology* **404**:241–258. DOI: https://doi.org/10.1113/jphysiol.1988. sp017288, PMID: 2855348

Llinás R, Yarom Y. 1986. Oscillatory properties of guinea-pig inferior olivary neurones and their pharmacological modulation: an in vitro study. *The Journal of Physiology* **376**:163–182. DOI: https://doi.org/10.1113/jphysiol. 1986.sp016147, PMID: 3795074

Loewenstein Y, Mahon S, Chadderton P, Kitamura K, Sompolinsky H, Yarom Y, Häusser M. 2005. Bistability of cerebellar purkinje cells modulated by sensory stimulation. *Nature Neuroscience* **8**:202–211. DOI: https://doi. org/10.1038/nn1393, PMID: 15665875

Marcaggi P, Attwell D. 2007. Short- and long-term depression of rat cerebellar parallel fibre synaptic transmission mediated by synaptic crosstalk. *The Journal of Physiology* **578**:545–550. DOI: https://doi.org/10. 1113/jphysiol.2006.115014, PMID: 17110417

Marr D. 1969. A theory of cerebellar cortex. *The Journal of Physiology* **202**:437–470. DOI: https://doi.org/10. 1113/jphysiol.1969.sp008820, PMID: 5784296

Maruta J, Hensbroek RA, Simpson JI. 2007. Intraburst and interburst signaling by climbing fibers. *Journal of Neuroscience* **27**:11263–11270. DOI: https://doi.org/10.1523/JNEUROSCI.2559-07.2007, PMID: 17942720

Mathy A, Ho SS, Davie JT, Duguid IC, Clark BA, Häusser M. 2009. Encoding of oscillations by axonal bursts in inferior olive neurons. *Neuron* **62**:388–399. DOI: https://doi.org/10.1016/j.neuron.2009.03.023, PMID: 19447094

McCormick DA, Clark GA, Lavond DG, Thompson RF. 1982. Initial localization of the memory trace for a basic form of learning. *PNAS* **79**:2731–2735. DOI: https://doi.org/10.1073/pnas.79.8.2731, PMID: 6953427

Medina JF, Lisberger SG. 2008. Links from complex spikes to local plasticity and motor learning in the cerebellum of awake-behaving monkeys. *Nature Neuroscience* **11**:1185–1192. DOI: https://doi.org/10.1038/nn. 2197, PMID: 18806784

Menzies JR, Porrill J, Dutia M, Dean P. 2010. Synaptic plasticity in medial vestibular nucleus neurons: comparison with computational requirements of VOR adaptation. *PLOS ONE* **5**:e13182. DOI: https://doi.org/10.1371/ journal.pone.0013182, PMID: 20957149

Minsky M. 1954. Theory of neural-analog reinforcement systems and its application to the brain-model problem. Princeton.

Minsky M. 1961. Steps toward artificial intelligence. *Proceedings of the IRE* **49**: 8–30. DOI: https://doi.org/10. 1109/JRPROC.1961.287775

Mittmann W, Häusser M. 2007. Linking synaptic plasticity and spike output at excitatory and inhibitory synapses onto cerebellar purkinje cells. *Journal of Neuroscience* **27**:5559–5570. DOI: https://doi.org/10.1523/ JNEUROSCI.5117-06.2007, PMID: 17522301

Mlonyeni M. 1973. The number of purkinje cells and inferior olivary neurones in the cat. *The Journal of Comparative Neurology* **147**:1–9. DOI: https://doi.org/10.1002/cne.901470102, PMID: 4682183

Monsivais P, Clark BA, Roth A, Häusser M. 2005. Determinants of action potential propagation in cerebellar purkinje cell axons. *Journal of Neuroscience* **25**:464–472. DOI: https://doi.org/10.1523/JNEUROSCI.3871-04. 2005, PMID: 15647490

Mooney R. 2009. Neurobiology of song learning. *Current Opinion in Neurobiology* **19**:654–660. DOI: https://doi. org/10.1016/j.conb.2009.10.004, PMID: 19892546

Mortimer JA. 1975. Cerebellar responses to teleceptive stimuli in alert monkeys. *Brain Research* **83**:369–390. DOI: https://doi.org/10.1016/0006-8993(75)90831-8, PMID: 163116

Mostofi A, Holtzman T, Grout AS, Yeo CH, Edgley SA. 2010. Electrophysiological localization of eyeblink-related microzones in rabbit cerebellar cortex. *Journal of Neuroscience* **30**:8920–8934. DOI: https://doi.org/10.1523/JNEUROSCI.6117-09.2010, PMID: 20592214

Murakami I. 2004. Correlations between fixation stability and visual motion sensitivity. *Vision Research* **44**:751–761. DOI: https://doi.org/10.1016/j.visres.2003.11.012, PMID: 14967202

Napper RM, Harvey RJ. 1988. Number of parallel fiber synapses on an individual purkinje cell in the cerebellum of the rat. *The Journal of Comparative Neurology* **274**:168–177. DOI: https://doi.org/10.1002/cne.902740204, PMID: 3209740

Nevian T, Sakmann B. 2006. Spine Ca2+ signaling in spike-timing-dependent plasticity. *Journal of Neuroscience* **26**:11001–11013. DOI: https://doi.org/10.1523/JNEUROSCI.1749-06.2006, PMID: 17065442

Nicholson C, ten Bruggencate G, Stöckle H, Steinberg R. 1978. Calcium and potassium changes in extracellular microenvironment of cat cerebellar cortex. *Journal of Neurophysiology* **41**:1026–1039. DOI: https://doi.org/10.1152/jn.1978.41.4.1026, PMID: 681986

Ohmae S, Medina JF. 2015. Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice. *Nature Neuroscience* **18**:1798–1803. DOI: https://doi.org/10.1038/nn.4167, PMID: 26551541

Ohyama T, Nores WL, Medina JF, Riusech FA, Mauk MD. 2006. Learning-induced plasticity in deep cerebellar nucleus. *Journal of Neuroscience* **26**:12656–12663. DOI: https://doi.org/10.1523/JNEUROSCI.4023-06.2006, PMID: 17151268

Olveczky BP, Andalman AS, Fee MS. 2005. Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLOS Biology* **3**:e153. DOI: https://doi.org/10.1371/journal.pbio.0030153, PMID: 15826219

Optican LM, Robinson DA. 1980. Cerebellar-dependent adaptive control of primate saccadic system. *Journal of Neurophysiology* **44**:1058–1076. DOI: https://doi.org/10.1152/jn.1980.44.6.1058, PMID: 7452323

Ozden I, Dombeck DA, Hoogland TM, Tank DW, Wang SS. 2012. Widespread state-dependent shifts in cerebellar activity in locomoting mice. *PLOS ONE* **7**:e42650. DOI: https://doi.org/10.1371/journal.pone.0042650, PMID: 22880068

Paoletti P, Ascher P, Neyton J. 1997. High-affinity zinc inhibition of NMDA NR1-NR2A receptors. *The Journal of Neuroscience* **17**:5711–5725. DOI: https://doi.org/10.1523/JNEUROSCI.17-15-05711.1997, PMID: 9221770

Pugh JR, Raman IM. 2006. Potentiation of mossy fiber EPSCs in the cerebellar nuclei by NMDA receptor activation followed by postinhibitory rebound current. *Neuron* **51**:113–123. DOI: https://doi.org/10.1016/j.neuron.2006.05.021, PMID: 16815336

Pugh JR, Raman IM. 2008. Mechanisms of potentiation of mossy fiber EPSCs in the cerebellar nuclei by coincident synaptic excitation and inhibition. *Journal of Neuroscience* **28**:10549–10560. DOI: https://doi.org/10.1523/JNEUROSCI.2061-08.2008, PMID: 18923031

R Core Team. 2013. R: A Language and Environment for Statistical Computing. Vienna, Austria, R Foundation for Statistical Computing. http://www.R-project.org/

Rasmussen A, Jirenhed DA, Zucca R, Johansson F, Svensson P, Hesslow G. 2013. Number of spikes in climbing fibers determines the direction of cerebellar learning. *Journal of Neuroscience* **33**:13436–13440. DOI: https://doi.org/10.1523/JNEUROSCI.1527-13.2013, PMID: 23946401

Rasmussen A, Jirenhed DA, Wetmore DZ, Hesslow G. 2014. Changes in complex spike activity during classical conditioning. *Frontiers in Neural Circuits* **8**:90. DOI: https://doi.org/10.3389/fncir.2014.00090, PMID: 25140129

Robbins H, Monro S. 1951. A stochastic approximation method. *The Annals of Mathematical Statistics* **22**:400–407. DOI: https://doi.org/10.1214/aoms/1177729586

Robinson DA. 1976. Adaptive gain control of vestibuloocular reflex by the cerebellum. *Journal of Neurophysiology* **39**:954–969. DOI: https://doi.org/10.1152/jn.1976.39.5.954, PMID: 1086347

Rogers RF, Fender AF, Steinmetz JE. 1999. The cerebellum is necessary for rabbit classical eyeblink conditioning with a non-somatosensory (photic) unconditioned stimulus. *Behavioural Brain Research* **104**:105–112. DOI: https://doi.org/10.1016/S0166-4328(99)00054-6, PMID: 11125728

Rowan MJM, Bonnan A, Zhang K, Amat SB, Kikuchi C, Taniguchi H, Augustine GJ, Christie JM. 2018. Graded Control of Climbing-Fiber-Mediated Plasticity and Learning by Inhibition in the Cerebellum. *Neuron* **99**:999–1015. DOI: https://doi.org/10.1016/j.neuron.2018.07.024, PMID: 30122378

Ruigrok TJ, Pijpers A, Goedknegt-Sabel E, Coulon P. 2008. Multiple cerebellar zones are involved in the control of individual muscles: a retrograde transneuronal tracing study with rabies virus in the rat. *European Journal of Neuroscience* **28**:181–200. DOI: https://doi.org/10.1111/j.1460-9568.2008.06294.x, PMID: 18662342

Safo P, Regehr WG. 2008. Timing dependence of the induction of cerebellar LTD. *Neuropharmacology* **54**:213–218. DOI: https://doi.org/10.1016/j.neuropharm.2007.05.029, PMID: 17669443

Sakurai M. 1987. Synaptic modification of parallel fibre-Purkinje cell transmission in in vitro guinea-pig cerebellar slices. *The Journal of Physiology* **394**:463–480. DOI: https://doi.org/10.1113/jphysiol.1987.sp016881, PMID: 2832595

Sarkisov DV, Wang SS. 2008. Order-dependent coincidence detection in cerebellar Purkinje neurons at the inositol trisphosphate receptor. *Journal of Neuroscience* **28**:133–142. DOI: https://doi.org/10.1523/JNEUROSCI.1729-07.2008, PMID: 18171931

Schild RF. 1970. On the inferior olive of the albino rat. *The Journal of Comparative Neurology* **140**:255–259. DOI: https://doi.org/10.1002/cne.901400302, PMID: 5476885

Schonewille M, Khosrovani S, Winkelman BH, Hoebeek FE, De Jeu MT, Larsen IM, Van der Burg J, Schmolesky MT, Frens MA, De Zeeuw CI. 2006a. Purkinje cells in awake behaving animals operate at the upstate

membrane potential. *Nature Neuroscience* **9**:459–461. DOI: https://doi.org/10.1038/nn0406-459, PMID: 1656 8098

Schonewille M, Luo C, Ruigrok TJ, Voogd J, Schmolesky MT, Rutteman M, Hoebeek FE, De Jeu MT, De Zeeuw CI. 2006b. Zonal organization of the mouse flocculus: physiology, input, and output. *The Journal of Comparative Neurology* **497**:670–682. DOI: https://doi.org/10.1002/cne.21036, PMID: 16739198

Schultz W. 1986. Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey. *Journal of Neurophysiology* **56**:1439–1461. DOI: https://doi.org/10.1152/jn.1986.56.5.1439, PMID: 3794777

Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. *Science* **275**:1593–1599. DOI: https://doi.org/10.1126/science.275.5306.1593, PMID: 9054347

Sejnowski TJ. 1977. Storing covariance with nonlinearly interacting neurons. *Journal of Mathematical Biology* **4**: 303–321. DOI: https://doi.org/10.1007/BF00275079, PMID: 925522

Seung HS. 2003. Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron* **40**:1063–1073. DOI: https://doi.org/10.1016/S0896-6273(03)00761-X, PMID: 14687542

Shalev-Shwartz S, Ben-David S. 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press . DOI: https://doi.org/10.1017/CBO9781107298019

Silver IA, Erecińska M. 1990. Intracellular and extracellular changes of [Ca2+] in hypoxia and ischemia in rat brain in vivo. *The Journal of General Physiology* **95**:837–866. DOI: https://doi.org/10.1085/jgp.95.5.837, PMID: 2163431

Soetedjo R, Kojima Y, Fuchs A. 2008. Complex spike activity signals the direction and size of dysmetric saccade errors. *Progress in brain research* **171**:153–159. DOI: https://doi.org/10.1016/S0079-6123(08)00620-1, PMID: 1 8718294

Stone LS, Lisberger SG. 1986. Detection of tracking errors by visual climbing fiber inputs to monkey cerebellar flocculus during pursuit eye movements. *Neuroscience Letters* **72**:163–168. DOI: https://doi.org/10.1016/0304-3940(86)90073-X, PMID: 3101005

Stuart G, Häusser M. 1994. Initiation and spread of sodium action potentials in cerebellar Purkinje cells. *Neuron* **13**:703–712. DOI: https://doi.org/10.1016/0896-6273(94)90037-X, PMID: 7917300

Sutton RS, Barto AG. 1998. *Introduction to Reinforcement Learning*. 1st ed. Cambridge, MA: MIT Press.

Suvrathan A, Payne HL, Raymond JL. 2016. Timing Rules for Synaptic Plasticity Matched to Behavioral Function. *Neuron* **92**:959–967. DOI: https://doi.org/10.1016/j.neuron.2016.10.022, PMID: 27839999

Suvrathan A, Raymond JL. 2018. Depressed by Learning-Heterogeneity of the Plasticity Rules at Parallel Fiber Synapses onto Purkinje Cells. *The Cerebellum*. DOI: https://doi.org/10.1007/s12311-018-0968-8, PMID: 3006 9835

Thach WT. 1968. Discharge of Purkinje and cerebellar nuclear neurons during rapidly alternating arm movements in the monkey. *Journal of Neurophysiology* **31**:785–797. DOI: https://doi.org/10.1152/jn.1968.31.5.785, PMID: 4974877

Wang SS, Denk W, Häusser M. 2000. Coincidence detection in single dendritic spines mediated by calcium release. *Nature Neuroscience* **3**:1266–1273. DOI: https://doi.org/10.1038/81792, PMID: 11100147

Wetmore DZ, Jirenhed DA, Rasmussen A, Johansson F, Schnitzer MJ, Hesslow G. 2014. Bidirectional plasticity of Purkinje cells matches temporal features of learning. *Journal of Neuroscience* **34**:1731–1737. DOI: https://doi.org/10.1523/JNEUROSCI.2883-13.2014, PMID: 24478355

Wigström H, Gustafsson B. 1983a. Facilitated induction of hippocampal long-lasting potentiation during blockade of inhibition. *Nature* **301**:603–604. DOI: https://doi.org/10.1038/301603a0, PMID: 6298626

Wigström H, Gustafsson B. 1983b. Large long-lasting potentiation in the dentate gyrus in vitro during blockade of inhibition. *Brain Research* **275**:153–158. DOI: https://doi.org/10.1016/0006-8993(83)90428-6, PMID: 6313124

Williams RJ. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* **8**:229–256. DOI: https://doi.org/10.1007/BF00992696

Xie X, Seung HS. 2004. Learning in neural networks by reinforcement of irregular spiking. *Physical Review E* **69**: 041909. DOI: https://doi.org/10.1103/PhysRevE.69.041909, PMID: 15169045

Yang Y, Lisberger SG. 2014. Purkinje-cell plasticity and cerebellar motor learning are graded by complex-spike duration. *Nature* **510**:529–532. DOI: https://doi.org/10.1038/nature13282, PMID: 24814344

Yartsev MM, Givon-Mayo R, Maller M, Donchin O. 2009. Pausing purkinje cells in the cerebellum of the awake cat. *Frontiers in Systems Neuroscience* **3**:2. DOI: https://doi.org/10.3389/neuro.06.002.2009, PMID: 19390639

Yeo CH, Hardiman MJ, Glickstein M. 1984. Discrete lesions of the cerebellar cortex abolish the classically conditioned nictitating membrane response of the rabbit. *Behavioural Brain Research* **13**:261–266. DOI: https://doi.org/10.1016/0166-4328(84)90168-2, PMID: 6508904

Yeo CH, Hesslow G. 1998. Cerebellum and conditioned reflexes. *Trends in Cognitive Sciences* **2**:322–330. DOI: https://doi.org/10.1016/S1364-6613(98)01219-4, PMID: 21227228

Zhang W, Linden DJ. 2006. Long-term depression at the mossy fiber-deep cerebellar nucleus synapse. *Journal of Neuroscience* **26**:6935–6944. DOI: https://doi.org/10.1523/JNEUROSCI.0784-06.2006, PMID: 16807323

Zucca R, Rasmussen A, Bengtsson F. 2016. Climbing Fiber Regulation of Spontaneous Purkinje Cell Activity and Cerebellum-Dependent Blink Responses(1,2,3). *eNeuro* **3**. :6. DOI: https://doi.org/10.1523/ENEURO.0067-15. 2015, PMID: 26839917

Appendix 1

Mathematical appendix

Reduced model

The model focuses on the rate $P(\tau)$ of one P-cell. Learning a pattern (movement) requires adjustment of the $P(\tau)$ rate to the target value R. The movement error is defined as

$$\mathcal{E}(\tau) = |P(\tau) - R| \tag{16}$$

Note that here τ represents trial/learning step number rather than temporal variations within a single movement realisation. In this simplified model, the P-cell activity during a movement is characterised by a single value rather than a sequence of values in time.

The other variable of the model besides $P(\tau)$ is $\mathcal{J}(\tau)$. It is used to quantify the strength

$$\mathcal{I}(\tau) = [\mathcal{J}(\tau) - qP(\tau)]_+ \tag{17}$$

of the $(NO \rightarrow \mathcal{I}O)$ inhibitory input which provides the current value of the estimated global error. The two terms in the expression of the current are meant to represent the excitatory action of the MF inputs on the discharge of the NO ($\mathcal{J}(\tau)$) and the inhibitory effects of the Purkinje cell discharge ($-qP(\tau)$). Learning steps consist of updates of the values $P(\tau)$ and $\mathcal{J}(\tau)$. They are of two kinds depending on whether the P-cell rate is perturbed or not.

Updates of type $\Delta P \Delta \mathcal{J}$ proceed as follows. The firing rate of the P-cell is increased by $A > 0$ and becomes $P(\tau) + A$. The value of the global error corresponding to this perturbed firing rate is thus

$$\mathcal{E}_p = |P(\tau) + A - R| \tag{18}$$

$P(\tau)$ and $\mathcal{J}(\tau)$ are updated depending on whether $\mathcal{E}_p$ is smaller or greater than the current value of the estimated global error $[\mathcal{J}(\tau) - qP(\tau) - qA]_+$. The parameter $q > 0$ quantifies how much the NO output is affected by the P-cell discharge, with the brackets indicating rectification (which imposes the constraint that the NO output remains non-negative).

- If $\mathcal{E}_p > [\mathcal{J}(\tau) - qP(\tau) - qA]_+$ the perturbation is judged to have increased the error and therefore to have the wrong sign: the perturbed firing rate needs to be decreased. Concurrently, $\mathcal{J}$ needs to be increased since $\mathcal{I}$ is judged too low compared to the real value of the error. Thus $P(\tau)$ and $\mathcal{J}(\tau)$ are changed to

$$P(\tau + 1) = [P(\tau) - \Delta P]_+ \tag{19}$$

$$\mathcal{J}(\tau + 1) = \mathcal{J}(\tau) + \Delta \mathcal{J} \tag{20}$$

As above the brackets indicate rectification, which imposes the constraint that firing rates are non-negative.

- If $\mathcal{E}_p < [\mathcal{J}(\tau) - qP(\tau) - qA]_+$, the converse reasoning leads to changes of $P(\tau)$ and $\mathcal{J}(\tau)$ in the opposite directions

$$P(\tau + 1) = P(\tau) + \Delta P \tag{21}$$

$$\mathcal{J}(\tau + 1) = [\mathcal{J}(\tau) - \Delta \mathcal{J}]_+ \tag{22}$$

Updates of type $\Delta\mathcal{J}$ are performed as described above but without any perturbation ($A = 0$) and without any update of P (only the estimated error $\mathcal{I}(\tau)$ is updated in updates of type $\Delta\mathcal{J}$). Namely,

- If $\mathcal{E} > [\mathcal{J}(\tau) - qP(\tau)]_+$ the estimated error is judged too low compared to the real value of the error and $\mathcal{J}$ needs to be increased. Thus $P(\tau)$ is kept unchanged and $\mathcal{J}(\tau)$ is modified

$$P(\tau+1) = P(\tau) \qquad (23)$$

$$\mathcal{J}(\tau+1) = \mathcal{J}(\tau) + \Delta\mathcal{J} \qquad (24)$$

- If $\mathcal{E} < [\mathcal{J}(\tau) - qP(\tau)]_+$, the converse reasoning leads to changes $\mathcal{J}(\tau)$ in the opposite directions

$$P(\tau+1) = P(\tau) \qquad (25)$$

$$\mathcal{J}(\tau+1) = [\mathcal{J}(\tau) - \Delta\mathcal{J}]_+ \qquad (26)$$

The proposed core operation of the algorithm for one movement can thus be described as follows. At each time step, a perturbation occurs with probability ρ. The occurrence of a perturbation gives rise to an update of type $\Delta P\Delta\mathcal{J}$, performed as described by **Equations 19-22**. For the complementary $(1 - \rho)$ fraction of time steps, no perturbation occurs ($A = 0$) and updates of type $\Delta\mathcal{J}$ are performed, as described by **Equations 23-26**.

This abstract model depends on five parameters: A, the amplitude of the rate perturbation, ΔP and $\Delta\mathcal{J}$, the update amplitudes of the rate and error estimate, respectively, q the strength of inhibition of the NO by the P-cell, and ρ, which describes the probability of an update of type $\Delta P\Delta\mathcal{J}$ (and the complementary probability $1 - \rho$ of $\Delta\mathcal{J}$ updates). The ratio of the update amplitudes play an important role and is denoted by β, $\beta = \Delta\mathcal{J}/\Delta P$. The target firing rates are chosen in the range $[0, R_{max}]$, which simply fixes the firing rate scale. We would like to determine the conditions on the five parameters $A, \Delta P, \Delta\mathcal{J}, q$ and ρ for the algorithm to 'converge'. For simplicity, we here consider only constant perturbation and learning steps. Therefore, at best the rate and the error fluctuates around the target rates and zero error, in a bounded domain of size determined by the magnitude of the constant perturbation and learning steps. We say that the algorithm converges when this situation is reached. We would also like to understand how they determine the rate of convergence and the residual error.

Convergence in the one-cell case

Stochastic gradient descent is conveniently analysed with a phase-plane description, by following the values $(P, \mathcal{J})$ from one update to the next in the $(P, \mathcal{J})$ plane. The dynamics randomly alternate between updates of type $\Delta P\Delta\mathcal{J}$ and $\Delta\mathcal{J}$, depending on whether the P-cell rate is perturbed by the CF or not. We consider these two types of update in turn.

For updates of type $\Delta P\Delta\mathcal{J}$, the update depends on whether the perturbed error $\mathcal{E}_p = |P - R + A|$ is larger or smaller than the estimated error $[\mathcal{J} - qP]_+$. Namely, it depends on the location of the current $(P, \mathcal{J})$ with respect to the two lines in the P-$\mathcal{J}$ plane, $C_\pm$, $\{(q \pm 1)P + (q \pm 1)A \mp R = \mathcal{J}\}$, of slopes $q \pm 1$ (see **Figure 9B**). The dynamics of **Equations 19-22** are such that each update moves the point $(P, \mathcal{J})$ by adding to it the vectorial increment $\pm(\Delta P, -\Delta\mathcal{J})$ with the $+$ sign holding in the quadrant above the lines $C_\pm$ and the minus sign elsewhere. These updates move the point $(P, \mathcal{J})$ towards the lines $C_\pm$. In the triangular domain below the line C_-, the update does not directly move the $(P, \mathcal{J})$ trajectory towards the C_- line, when $\beta < (1 - q)$ and the update has an angle greater than the inclination of the C_-. However, it reduces P. When $P = 0$ below C_-, the updates are strictly upward and towards C_-.

Updates of type $\Delta\mathcal{J}$ depend on the location of the point $(P, \mathcal{J})$ with respect to the lines $D_\pm$, $\{(q \pm 1)P \mp R = \mathcal{J}\}$, since the cell firing rate is not perturbed in these updates. In the quadrant above the lines $D_\pm$, an update moves the point $(P, \mathcal{J})$ downwards by $\Delta\mathcal{J}$,

in other words, adds the vectorial increment $(0, -\Delta\mathcal{J})$. In the complementary domain of the $(P, \mathcal{J})$ plane, an update moves the point $(P, \mathcal{J})$ upwards by $\Delta\mathcal{J}$, in other words, adds the opposite vectorial increment $(0, \Delta\mathcal{J})$. Both updates move $(P, \mathcal{J})$ towards the lines $D_\pm$.

Learning proceeds by performing updates of type $\Delta P\Delta\mathcal{J}$ and $\Delta\mathcal{J}$ with respective probabilities ρ and $(1 - \rho)$. Starting from an initial coordinate $(P, \mathcal{J})$, three phases can be distinguished as described in the main text.

Above the lines C_+ and D_-, the mixed updates lead the point $(P, \mathcal{J})$ to perform a random walk with a systematic mean rightward-downward drift per update equal to

$$\rho(\Delta P, -\Delta\mathcal{J}) + (1 - \rho)(0, -\Delta\mathcal{J}) = (\rho\Delta P, -\Delta\mathcal{J}), \tag{27}$$

for $(P, \mathcal{J})$ above C_+ and D_-. Below the lines C_- and D_+, the updates are opposite and the mean drift per update is leftward-upward, equal to

$$(-\rho\Delta P, \Delta\mathcal{J}), \tag{28}$$

for $(P, \mathcal{J})$ below C_- and D_+. Depending on its initial condition and the exact set of updates drawn, this leads $(P, \mathcal{J})$ to reach either one of the two 'convergence corridors', between the lines C_+ and D_+, or between the lines C_- and D_- (see *Figure 9B*). In the triangular domain below the line C_-, the mean leftward-upward drift has an angle greater than the inclination of the C_- line for $\rho\Delta P(1 - q) > \Delta\mathcal{J}$ (i.e. $\beta < (1 - q)\rho$). In this case, the mean drift does not ensure that the $(P, \mathcal{J})$ trajectory crosses the C_- line. However, if P becomes zero before crossing C_-, the positivity constraint on the rate, imposes that subsequent updates are strictly upward, until $(P, \mathcal{J})$ reaches C_-. Examples of such trajectories can be seen in *Appendix 1—figure 1C and D* (red lines).
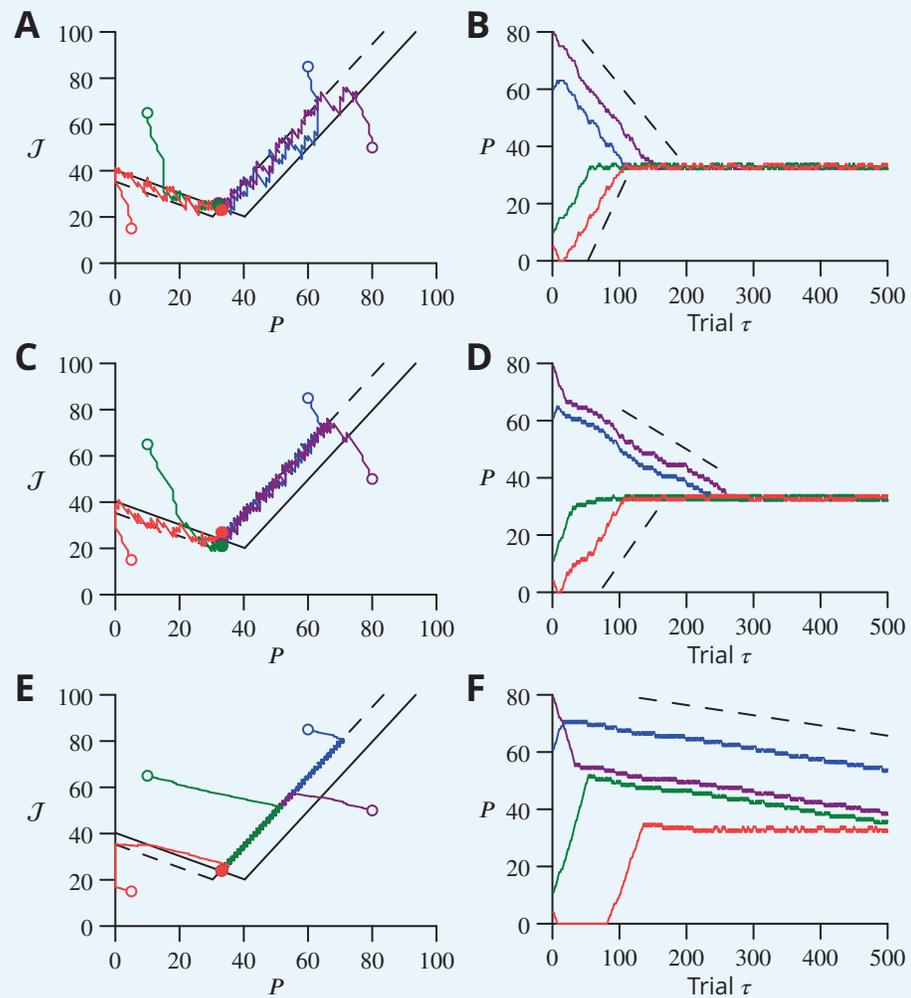
Appendix 1—figure 1. Reduced model: different cases of stochastic gradient learning are illustrated, for the parameters marked by *solid red circles* in **Figure 9C**, except for the case already shown in **Figure 9E and F**. (**A**) Dynamics in the $(P, \mathcal{J})$ plane for $\rho = 0.5$ and $\beta = 2$ ($A = 10, \Delta P = 1, \Delta \mathcal{J} = 2$). Trajectories (*solid lines*) from different initial conditions (*open circles*) are represented in the $(P, \mathcal{J})$ plane. The trajectories converge by oscillating around C_+ in the C_+/D_+ 'corridor' and around D_- in the C_-/D_- 'corridor'. Trajectory endings are marked by *filled circles*. (**B**) Time courses of $P(\tau)$ with corresponding colours. The slope of convergence (*dotted lines*) predicted by **Equation 32 and 33** agrees well with the observed convergence rate for C_+ while it is less accurate for D_- when the assumption that the 'corridor' width is much greater than ΔP and $\Delta \mathcal{J}$ does not hold. (**C, D**) Same graphs for $\rho = 0.75$. The trajectories converge by oscillating around C_+ in the C_+/D_+ and around C_- in the C_-/D_- 'corridor'. (**E, F**) Same graphs for $\rho = 0.75$ and $\beta = 0.25$ (i.e. $A = 2, \Delta P = 1, \Delta \mathcal{J} = 0.25$). The trajectories converge by oscillating around C_+, which is the only attractive line.

DOI: https://doi.org/10.7554/eLife.31599.019

When a corridor is reached, in a second phase, $(P, \mathcal{J})$ follows a stochastic walk in the corridor.

A sufficient condition for convergence is that a single update cannot cross the two boundary lines of a corridor at once. For the (C_+, D_+) corridor, the crossing by a single $\Delta P \Delta \mathcal{J}$ update provides the most stringent requirement,.

$$A > \Delta \mathcal{J}/(q+1) + \Delta P = \Delta \mathcal{J}[1/(q+1) + 1/\beta]. \tag{29}$$

The crossing of the (C_-, D_-) corridor by a single $\Delta \mathcal{J}$ provides the other requirement

$$A(1-q) > \Delta\mathcal{J} \tag{30}$$

When these conditions are met, alternation between the two types of updates produces a mean downward drift of the error. This downward drift controls the convergence rate and depends on the relative size of the perturbation and the discrete ΔP and $\Delta\mathcal{J}$ modifications (in other words, the number of modifications that are needed to cross the convergence corridor). For simplicity, we consider the case where the perturbation A is large compared to ΔP and $\Delta\mathcal{J}$, so that oscillations basically take place around one line of the corridor, as illustrated in **Figure 9E** and **Appendix 1—figure 1**. This amounts to being able to neglect the probability of crossing the other line of the corridor. The convergence rate then depends on the corridor line around which it takes place and can be obtained by noting that around a given corridor line, one of the two types of updates always has the same sign.

In the (C_+, D_+) corridor, performing type $\Delta P \Delta\mathcal{J}$ updates for fraction ρ of the steps and type $\Delta\mathcal{J}$ updates for the complementary fraction $(1-\rho)$ leads to the average displacement per step,

$$\rho(-\Delta P, \Delta\mathcal{J}) + (1-\rho)(0, -\Delta\mathcal{J}) = (-\rho\Delta P, (2\rho-1)\Delta\mathcal{J}) \tag{31}$$

Thus, as summarised in **Figure 9C**

- For $\rho < \beta/[(q+1)+2\beta]$ the average displacement leads to D_+ (**Figure 9E**). The average downward drift in the corridor can be obtained by noting that for a large A, the $(P, \mathcal{J})$ trajectory has a negligibly small probability of crossing the C_+ line. Therefore, all the type $\Delta P \Delta\mathcal{J}$ updates are of the same sign, of the form $(-\Delta P, \Delta\mathcal{J})$, and chosen with probability ρ. Since updates of type $\Delta\mathcal{J}$ do not change the value of P, P approaches its target rate with a mean speed per step V_c,

$$D_+: V_c = -\rho\Delta P \tag{32}$$

A comparison between this computed drift and simulated trajectories is shown in **Figure 9F**.

- For $\rho > \beta/[(q+1)+2\beta]$, the average displacement leads to C_+ (**Appendix 1—figure 1A,C,E**). In this case, all type $\Delta\mathcal{J}$ updates are of the same sign, of the form $(0, -\Delta\mathcal{J})$. In contrast, a fraction f of type $\Delta P \Delta\mathcal{J}$ updates are of the form $(-\Delta P, \Delta\mathcal{J})$, while a fraction $(1-f)$ is of the opposite form $(\Delta P, -\Delta\mathcal{J})$ with f to be determined. The average drift is thus $(1-\rho)(0, -\Delta\mathcal{J}) + \rho[f(-\Delta P, \Delta\mathcal{J}) + (1-f)(\Delta P, -\Delta\mathcal{J})] = (\rho(1-2f)\Delta P, [(\rho-1) + \rho(2f-1)]\Delta\mathcal{J})$. Requiring that this drift has slope $(1+q)$, like C_+, gives $2f-1 = \beta(1-\rho)/[(1+q+\beta)\rho]$ and $f = [\beta + \rho(1+q)]/[2\rho(1+q+\beta)]$ (which indeed obeys $0 < f < 1$ in the parameter domain considered). Thus, P approaches its target rate with a mean speed per step V_c,

$$C_+: V_c = -\frac{1-\rho}{1+q+\beta}\Delta\mathcal{J} \tag{33}$$

A comparison between this computed drift and simulated trajectories is shown in **Appendix 1—figure 1B,D,F**.

In the (C_-, D_-) corridor, the average drift per step is opposite to the drift in the (C_+, D_+) corridor (**Equation 31**). Thus,

- For $\beta > (1-q)$ and $\rho > \beta/(2\beta + q - 1)$ the average displacement leads to C_- (**Appendix 1—figure 1A**). Updates of type $\Delta\mathcal{J}$ are always of the form $(0, \Delta\mathcal{J})$, while a fraction f of type $\Delta P \Delta\mathcal{J}$ updates are of the form $(\Delta P, -\Delta\mathcal{J})$ and a fraction $(1-f)$ are of the opposite form $(-\Delta P, \Delta\mathcal{J})$. Again, since the average drift is along C_- of slope $q-1$, one obtains $2f-1 = \beta(1-\rho)/[(\beta+q-1)\rho]$ or $f = [\beta + \rho(q-1)]/[2\rho(\beta + q - 1)]$ (which obeys $0 < f < 1$ in the parameter domain considered). The convergence speed V_c is thus,

$$C_-: V_c = \frac{1-\rho}{\beta + q - 1} \Delta \mathcal{J} \tag{34}$$

A comparison between this computed drift and simulated trajectories is shown in **Appendix 1—figure 1B**.

- For $\beta < 1 - q$, or $\{\beta > 1, \rho < \beta(2\beta - 1)\}$, the complementary parameter domain, the drift in the corridor leads to D_- (**Appendix 1—figure 1A,C**). The domain $\beta < (1 - q)\rho$ can be excluded, since when the point $(P, \mathcal{J})$ crosses D_-, the drift in the upper quadrant (C_+, D_-) (**Equation 27**) tends to bring it to the other corridor (C_+, D_+) (**Appendix 1—figure 1E**). Near the line D_-, type $\Delta P \Delta \mathcal{J}$ updates are always of the form $(\Delta P, -\Delta \mathcal{J})$. The convergence speed V_c is thus

$$D_-: V_c = \rho \Delta P \tag{35}$$

A comparison between this computed drift and simulated trajectories is shown in **Figure 9F** and **Appendix 1—figure 1B**.

Perceptron model and simulations

The architecture of the perceptron model is again the simplified circuit of **Figure 9A**, with $N_M = 1000$ mossy fibres projecting onto a single P-cell, with weights w_i, $i = 1, \ldots, N_M$ (which are all positive or zero). N_P patterns are generated randomly. Activities of mossy fibres in different patterns are i.i.d. binary random variables M_i^μ with coding level f (i.e. $M_i = 1$ with probability f and 0 with probability $1 - f$; in the present simulations $f = 0.2$). The P-cell desired rates for the N_P patterns are i.i.d. uniform variables R^μ from 0 to P_{max} (=100 Hz).

In the case of the stochastic gradient descent with estimated global errors algorithm, at each trial a pattern μ is randomly drawn without replacement among the total N_P and there is a probability ρ of a perturbation of amplitude A. The output of the P-cell is

$$P^\mu = \left[\frac{1}{\sqrt{N_M}} \left(\sum_i w_i M_i^\mu - \theta N_M \right) \right]_+ + \eta A \tag{36}$$

where $\eta = 1$ with probability ρ and $\eta = 0$ otherwise, which thus introduces a random perturbation of amplitude A into the P-cell firing. The error is defined as $\mathcal{E} = |P^\mu - R^\mu|$. Comparison with previously obtained results for the capacity of this analog perceptron (**Clopath and Brunel, 2013**) motivates our choice of weights normalisation and the parameterisation of the threshold as $\theta = \frac{1}{2} P_{max} \left(\sqrt{\frac{f}{3(1-f)\gamma}} - \frac{1}{\sqrt{N_M}} \right)$, where γ is a composite parameter reflecting the statistics of input and output firing, but here is equal to one.

An inferior olivary neurone receives the excitatory error signal but it also receives inhibitory inputs from the nucleo-olivary neurones driven by the mossy fibre inputs (which we have denoted M above), with weights v_i. These are also plastic and sign-constrained. They represent the current estimated error. The net drive of the inferior olivary neurone is

$$IO = \mathcal{E} - \mathcal{I}, \text{ with } \mathcal{I} = \left[\frac{1}{\sqrt{N_M}} \left(\sum_i v_i M_i^\mu - \theta N_M \right) - q P^\mu \right]_+ . \tag{37}$$

Here, the term qP^μ represents the specific inhibition of nucleo-olivary neurones by Purkinje cells and the term proportional to θ represents non-specific inhibition. In other words, $\frac{1}{\sqrt{N_M}} \left(\sum_i v_i M_i^\mu - \theta N_M \right)$ corresponds to $\mathcal{J}$ in **Equation 17**. For simplicity we assume the non-specific inhibition to be constant and the value of θ equal to that used in **Equation 36**.

The climbing fibre signal controlling plasticity is

$$c = \text{sign}(IO)$$

Weights are changed according to the following rule. Weights of $M \to P$ synapses active simultaneously with a perturbation are increased if c is negative and decreased if it is positive. Weights of $M \to NO$ synapses are increased if c is positive and decreased if not. Thus

$$w_i = \left[w_i - \alpha_w c \eta M_i^{\mu} \right]_+ \tag{38}$$

$$v_i = \left[v_i + \alpha_v c M_i^{\mu} \right]_+ \tag{39}$$

with the brackets indicating rectification (to impose the excitatory constraint).

The parameters of the simplified model of the previous subsection can be written as a function of those controlling the learning process in the present analog perceptron. The probability ρ of the two types of updates (with and without perturbation) and the amplitude of perturbation A are clearly identical in both models. In order to relate the previous ΔP and $\Delta \mathcal{J}$ to the present amplitude change of the weights α_w at mossy fibre–P-cell inputs and α_v for the indirect drive to the inferior olive from mossy fibres, we neglect the rectification constraints in *Equations 38 and 39* which is valid as long as synaptic weights are not very small (i.e. comparable to α_w or α_v; see below for further discussion). Therefore, the weight modifications result in the changes ΔP, of the perceptron firing rate, and $\Delta \mathcal{J}$, of the inhibitory input $\mathcal{I}$ to the olive,

$$\Delta P = \frac{1}{\sqrt{N_M}} \sum_i \alpha_w M_i^{\mu} = \alpha_w f \sqrt{N_M} \tag{40}$$

$$\Delta \mathcal{J} = \frac{1}{\sqrt{N_M}} \sum_i \alpha_v M_i^{\mu} = \alpha_v f \sqrt{N_M} \tag{41}$$

The convergence rate estimates of the previous section are compared with direct stochastic gradient descent learning simulations of the analog perceptrons in *Figure 10*. As shown in *Figure 10A*, for many single patterns, the convergence rate agrees well with the estimate *Equation 33*. However, for a fraction of the trajectories, the rate of convergence is about half of the prediction. This arises because these simulations constrained synaptic weights to be non-negative, which creates a significant fraction of synapses with negligible weights (*Brunel et al., 2004*) and produces a smaller effective step ΔP than estimated without taking this positivity constraint into account. For a larger number of patterns below the maximal learning capacity, the convergence rate per pattern is slower by a factor of $\approx 1.5 - 5$ for N_P between $100 - 350$ (*Figure 10B*).

The SGDEGE algorithm is compared to the usual delta rule in *Figure 10*. For the delta rule, the patterns are presented sequentially. When the pattern μ is presented, the weights are changed according to the signed error $\mathcal{E}_s$ as,

$$w_i = \left[w_i - \alpha_w M_i^{\mu} \mathcal{E}_s \right]_+ \tag{42}$$

The error is defined by the distance (positive or negative) of the P-cell firing rate to the target rate, $\mathcal{E}_s = P^{\mu} - R^{\mu}$, and P^{μ} is given as before by *Equation 36* (with $A = 0$).

The convergence of the SGDEGE algorithm is considerably slower than that obtained for the delta rule (*Figure 10C*), while the relative slowing due to interference of multiple patterns is comparable for the delta rule and the proposed algorithm (compare panels B,C in *Figure 10*). The slower convergence rate of the SGDEGE algorithm arises from the use of update steps of fixed amplitude while the update is proportional to the error magnitude for the delta rule. A modified delta rule with constant amplitude updates would give a convergence rate comparable to the SGDEGE one.