

Lectures 9 - 10: Beyond Brownian motion

Although Brownian motion is extremely useful in many contexts, it relies on assumptions that limit its applicability. In particular, the standard random walk (or Brownian) model assumes that jump events occur at *fixed* time intervals or, equivalently, that the waiting time between successive jumps is a constant. For a simple random walk in one dimension we take the time between steps to be some constant Δt . In Einstein's model of diffusion the characteristic time τ is identified with the mean collision time between particles. Both descriptions assume an average timescale but ignore the actual random nature of molecular collisions.

Further, a major assumption in Brownian motion is that we observe the dynamics only at the length (x) and time (t) scales much larger than the jump-length and collision timescales. Often, in diffusion within molecular systems, this assumption becomes difficult to satisfy, as the observed length scales can be comparable to the jump-lengths. A typical case of this is observed in QENS experiments on studying water diffusion.

9.1 Viability of Brownian assumptions

9.1.1 QENS Signatures of Water

A key experimental motivation for going beyond the Brownian model in liquid water comes from the quasielastic neutron scattering (QENS) measurements of Teixeira et al. (Phys. Rev. A 31, 1913 (1985)). Two central observations are relevant for our discussion: (i) the QENS spectra $S(Q, \omega)$ measured at different momentum transfers Q , and (ii) the dependence of the Lorentzian half-width at half-maximum (HWHM), $\Gamma(Q)$, on Q^2 . Both results provide direct evidence that water self-diffusion at molecular length scales does not follow the Brownian prediction $\Gamma(Q) = DQ^2$.

QENS line shapes at different Q

Figure 1 shows the experimental $S(Q, \omega)$ for water at $T = -5^\circ\text{C}$ for several values of Q . Each spectrum is Lorentzian in form, but its width increases strongly with Q . Specifically:

- At $Q = 1.05 \text{ \AA}^{-1}$, the line is narrow, indicating relatively slow decorrelation at this length scale.
- At $Q = 1.44 \text{ \AA}^{-1}$, the width is noticeably larger.
- At $Q = 2.03 \text{ \AA}^{-1}$, the line is significantly broadened, corresponding to sub-picosecond decorrelation times.

This strong, nonlinear broadening is inconsistent with Brownian diffusion, which predicts a width proportional to Q^2 . Instead, the behaviour indicates that water displacements at ångström scales arise from a small number of molecular jumps rather than from many small Brownian steps.




Figure 1: QENS spectra $S(Q, \omega)$ of water at $T = -5^\circ\text{C}$ from Teixeira et al. (1985). The Lorentzian width increases strongly with Q , deviating from the hydrodynamic (Brownian) prediction $\Gamma(Q) = DQ^2$ from Teixeira et al. (Phys. Rev. A 31, 1913 (1985)).

Variation of $\Gamma(Q)$ with Q^2

Figure 2 shows the dependence of the HWHM $\Gamma(Q)$ on Q^2 for temperatures between 20°C and -20°C . If diffusion were Brownian at these length scales, the data would follow straight lines:

$$\Gamma(Q) = DQ^2.$$

Instead, the curvature is pronounced:

- $\Gamma(Q)$ increases approximately linearly with Q^2 only at very small Q .
- At larger Q^2 , the growth saturates and becomes strongly subquadratic.
- Lower temperatures exhibit stronger deviation from linearity.

9.1.2 Collision time from kinetic theory

In order to understand why water and large solutes (e.g. proteins such as BSA) behave so differently in quasielastic neutron scattering (QENS), it is essential to quantify the *microscopic displacement produced by a single “kick”* from the solvent. A very simple and entirely classical route is obtained by using basic kinetic theory:

$$\tau_{\text{coll}} = \frac{1}{n \sigma v_{\text{th}}},$$

where n is the molecular number density, σ the collision cross section, and v_{th} the thermal speed. Throughout this derivation we will use Boltzmann’s constant k_B and absolute temperature T . The




Figure 2: Experimental variation of the QENS Lorentzian width $\Gamma(Q)$ versus Q^2 for water at different temperatures (Texeira et al. (Phys. Rev. A 31, 1913 (1985))). The strong curvature indicates breakdown of the Brownian scaling $\Gamma(Q) = DQ^2$.

goal of this section is to derive explicit expressions for these quantities and to compute the resulting microscopic displacement for both water molecules and large solutes such as proteins.

The number density of molecules in a fluid is related to its mass density ρ and molar mass M via Avogadro's number N_A :

$$n = \frac{\rho}{M} N_A.$$

For liquid water at room temperature the mass density is $\rho = 1000 \text{ kg/m}^3$ and the molar mass is $M = 18 \times 10^{-3} \text{ kg/mol}$. Substituting these numerical values and $N_A \approx 6.022 \times 10^{23} \text{ mol}^{-1}$ gives

$$n \approx \frac{1000}{18 \times 10^{-3}} (6.022 \times 10^{23}) \approx 3.3 \times 10^{28} \text{ m}^{-3}.$$

The collisional cross section can be approximated by treating a molecule as a hard sphere of effective diameter d . For a water molecule a reasonable estimate is $d \approx 3 \text{ \AA} = 3 \times 10^{-10} \text{ m}$. The cross-sectional area is then

$$\sigma = \pi d^2 \approx \pi (3 \times 10^{-10})^2 \approx 3 \times 10^{-19} \text{ m}^2.$$

A basic result of kinetic theory relates the thermal speed to the molecular mass m_w and temperature:

$$v_{\text{th}} = \sqrt{\frac{3k_B T}{m_w}}.$$

Here $m_w = M/N_A = 18 \times 10^{-3} \text{ kg/mol} / (6.022 \times 10^{23} \text{ mol}^{-1}) \approx 3.0 \times 10^{-26} \text{ kg}$. At $T = 300 \text{ K}$ the thermal speed is thus

$$v_{\text{th}} \approx \sqrt{\frac{3(1.38 \times 10^{-23} \text{ J/K})(300 \text{ K})}{3.0 \times 10^{-26} \text{ kg}}} \approx 600 \text{ m/s}.$$

Combining the values of n , σ and v_{th} , the collision time is

$$\tau_{\text{coll}} = \frac{1}{n \sigma v_{\text{th}}} = \frac{1}{(3.3 \times 10^{28})(3 \times 10^{-19})(600)} \approx 1.7 \times 10^{-13} \text{ s}.$$

This is the typical duration of a single microscopic impulse (or “kick”) imparted by a solvent molecule. Because this timescale depends solely on solvent parameters it is the same for water self-diffusion and for the centre-of-mass motion of any solute immersed in water, including large proteins.

9.1.2.1 Momentum transfer per kick

A single collision with a water molecule imparts an impulse of order the thermal momentum p_{th} . For a molecule of mass m , this momentum scale is

$$p_{\text{th}} \sim \sqrt{mk_B T}.$$

For a water molecule, $m = m_w = 3.0 \times 10^{-26} \text{ kg}$, giving

$$\Delta p \approx \sqrt{(3.0 \times 10^{-26})(1.38 \times 10^{-23})(300)} \approx 1.1 \times 10^{-23} \text{ kg m/s}.$$

Importantly, this momentum transfer scale is the same whether the impulse is delivered to another water molecule or to a large macromolecule such as BSA. The difference between water and BSA enters via the mass that receives this impulse.

9.1.2.2 Velocity jumps for water and BSA

The immediate velocity change due to a single impulse is

$$\Delta v = \frac{\Delta p}{m},$$

where m is the mass of the particle being kicked. For a water molecule of mass m_w the velocity jump is

$$\Delta v_w = \frac{1.1 \times 10^{-23}}{3.0 \times 10^{-26}} \approx 3.7 \times 10^2 \text{ m/s}.$$

In contrast, for a bovine serum albumin (BSA) molecule the centre-of-mass mass is roughly $m_{\text{BSA}} \approx 1.1 \times 10^{-22} \text{ kg}$. The same impulse produces a velocity increment

$$\Delta v_{\text{BSA}} = \frac{1.1 \times 10^{-23}}{1.1 \times 10^{-22}} \approx 1 \text{ m/s}.$$

Thus a solvent collision produces a velocity jump that is about 370 times larger for water than for the BSA centre-of-mass.

9.1.2.3 Spatial displacement per kick

The microscopic spatial displacement resulting from a single impulse combines the velocity increment Δv with the duration of the kick τ_{coll} :

$$\Delta = \Delta v \tau_{\text{coll}}.$$

Substituting Δv and τ_{coll} for water,

$$\Delta_{\text{w}} \approx (3.7 \times 10^2)(1.7 \times 10^{-13}) \approx 6.3 \times 10^{-11} \text{ m} = 0.63 \text{ \AA},$$

where we have used $1 \text{ \AA} = 10^{-10} \text{ m}$. For BSA the same impulse gives

$$\Delta_{\text{BSA}} \approx (1)(1.7 \times 10^{-13}) = 1.7 \times 10^{-13} \text{ m} = 1.7 \times 10^{-3} \text{ \AA}.$$

9.1.3 Implications for QENS

Quasielastic neutron scattering (QENS) experiments typically probe displacements on the order of $1\text{--}3 \text{ \AA}$, corresponding to scattering wavevectors $Q \sim 0.3\text{--}2 \text{ \AA}^{-1}$. The microscopic estimates above lead to the following qualitative conclusions:

- For water, a single kick displaces the molecule by approximately 0.6 \AA . Only a handful of statistically independent kicks are needed to reach the \AA -scale displacements probed in QENS. Therefore the displacement distribution for water self-diffusion need not be Gaussian, and the usual hydrodynamic scaling $\Gamma(Q) = DQ^2$ (where Γ is the relaxation rate and D the diffusion coefficient) does not apply across the entire QENS range.
- For a large macromolecule like BSA, one kick produces a displacement of only $\sim 10^{-3} \text{ \AA}$. To move by a few \AA the protein centre-of-mass must accumulate on the order of $10^3\text{--}10^4$ independent kicks. This places the protein in the Brownian (Gaussian displacement) regime over the QENS length scales, and the hydrodynamic scaling $\Gamma(Q) = DQ^2$ becomes valid.

Finally, note that both water and BSA experience kicks with Poisson-distributed inter-event times, so both yield Lorentzian line shapes in $S(Q, \omega)$. The difference arises because only BSA receives a sufficiently large number of spatially tiny kicks to produce Gaussian displacements on the \AA length scales relevant for QENS, whereas water does not. This microscopic picture provides a direct explanation of why Brownian diffusion works for macromolecules but breaks down for water self-diffusion at the experimental Q -range.

9.2 Continuous Time Random Walk

Consider molecular diffusion in liquids: the process is driven by random collisions with neighbouring particles or molecules. These collisions occur at random instants of time. Therefore the assumption of jumps at fixed time intervals is inadequate when describing the timing of molecular collisions. To go beyond this restriction we must *randomize the time intervals* between successive jump events and treat the waiting times as random variables drawn from some probability distribution.

The continuous time random walk (CTRW) model addresses this limitation by allowing both the waiting times between successive jumps and the jump lengths themselves to be random variables. Specifically, one introduces a waiting-time probability density function (PDF) $\psi(\tau)$ for the time interval between successive steps and a jump-length PDF $\lambda(\Delta x)$ for the spatial displacement during

a jump. By decoupling the distributions of waiting times and jump lengths, CTRW provides a more flexible framework for describing a wide range of physical, chemical and biological transport processes, particularly when:

- **Temporal randomness is essential.** Waiting times τ_i can be drawn from any appropriate distribution. For example, an exponential distribution represents Poissonian collision times. Heavy-tailed distributions (such as power laws) can model trapping or hindrance effects in disordered media.
- **Microscopic accuracy is desired.** For systems such as molecular liquids or colloidal suspensions, CTRW bridges microscopic kinetic models and macroscopic transport by incorporating both spatial and temporal heterogeneity.

In summary, CTRW offers a minimal yet powerful generalization of the random-walk framework by decoupling space and time. It is especially well-suited for systems where the timing of transitions is not regular, enabling it to describe anomalous transport phenomena more accurately than traditional random-walk models. In this chapter we will build the CTRW model from basic principles and explore its connections to Brownian motion, focusing on the physical implications for diffusion processes.

9.3 Construction and solution of CTRW

To construct a CTRW model we must randomize both the timing of jumps and their spatial lengths. We first focus on randomizing the time intervals between jumps and construct a stochastic *counting process* $N(t)$ that keeps track of how many events (jumps) have occurred up to time t . In the next subsection we will incorporate random jump lengths.

9.3.1 Modelling a stochastic counting process

Let $N(t)$ be a non-decreasing stochastic process that counts the number of events that have occurred by time t . The randomness arises because the events occur at random intervals. We assume that the waiting times τ_1, τ_2, \dots between successive events are independent and identically distributed (i.i.d.) random variables drawn from a waiting-time PDF $\phi(\tau)$. The survival function or no-event probability is

$$\Phi(t) = \int_t^\infty \phi(t') dt' = 1 - \int_0^t \phi(t') dt',$$

which gives the probability that no event has occurred by time t .

We are interested in the probability $\mathcal{P}(N, t)$ that exactly N events have occurred up to time t . We derive expressions for $\mathcal{P}(N, t)$ by elementary convolution arguments.

Probability of $N = 0$ events. The probability that no event occurs in the interval $(0, t)$ is simply the survival probability:

$$\mathcal{P}(0, t) = \Phi(t) = 1 - \int_0^t \phi(t') dt'. \quad (9.3.1)$$

Probability of $N = 1$ event. To have one event occur in $(0, t)$ we require that no event occurs in $(0, t')$ and then a single event occurs in (t', t) . The times t' at which the first event could occur are mutually exclusive and range from 0 to t . By independence of the waiting-time random variables, the probability is

$$\mathcal{P}(1, t) = \int_0^t \phi(t - t') \Phi(t') dt' = \phi \otimes \Phi(t),$$

where \otimes denotes convolution: $(f \otimes g)(t) = \int_0^t f(t - t')g(t') dt'$.

Probability of $N = 2$ events. To have two events in $(0, t)$ we split the interval into three subintervals $(0, t'')$, (t'', t') and (t', t) with $0 < t'' < t' < t$. We require that no event occurs in $(0, t'')$, exactly one event in (t'', t') , and exactly one event in (t', t) . By independence and mutual exclusivity of the time subdivisions,

$$\mathcal{P}(2, t) = \int_0^t \int_0^{t'} \phi(t - t') \phi(t' - t'') \Phi(t'') dt'' dt' = \phi \otimes \phi \otimes \Phi(t) = \phi^{\otimes 2} \otimes \Phi(t).$$

Case of general N . Iterating this argument yields a simple convolution structure:

$$\mathcal{P}(N, t) = \phi^{\otimes N}(t) \otimes \Phi(t). \quad (9.3.2)$$

That is, to have exactly N events by time t we must have N waiting times distributed according to ϕ , and then the final waiting time extends beyond t so that there is no $(N + 1)$ th event in $(0, t)$. Each convolution $\phi^{\otimes N}$ indicates N -fold convolution of ϕ with itself.

Laplace transform. It is convenient to take Laplace transforms of these probabilities. Define

$$\tilde{\mathcal{P}}(N, s) = \int_0^\infty e^{-st} \mathcal{P}(N, t) dt, \quad \tilde{\phi}(s) = \int_0^\infty e^{-st} \phi(t) dt, \quad \tilde{\Phi}(s) = \int_0^\infty e^{-st} \Phi(t) dt.$$

From (9.3.1) we find

$$\tilde{\Phi}(s) = \frac{1 - \tilde{\phi}(s)}{s}.$$

Because convolution in the time domain becomes multiplication in Laplace space, from (9.3.2) we obtain the general result

$$\tilde{\mathcal{P}}(N, s) = \left[\tilde{\phi}(s) \right]^N \frac{1 - \tilde{\phi}(s)}{s}. \quad (9.3.3)$$

9.3.1.1 Poisson Process: Exponential Waiting-Time Distribution

A particularly important case arises when the waiting-time PDF is exponential. Suppose

$$\phi(t) = \lambda e^{-\lambda t}, \quad t \geq 0,$$

where $\lambda > 0$ is the rate parameter and $\lambda^{-1} \equiv \tau_0$ is the mean waiting time. Differentiating ϕ reveals the memoryless (Markov) property:

$$\frac{d\phi(t)}{dt} = -\lambda\phi(t),$$

which in discrete form states that $\phi(t + \Delta t) \propto \phi(t)$. In other words, the probability distribution of the waiting time is independent of how much time has already passed since the last event.

The Laplace transform of the exponential PDF is

$$\tilde{\phi}(s) = \int_0^\infty e^{-st} \lambda e^{-\lambda t} dt = \frac{\lambda}{s + \lambda}.$$

Substituting into (9.3.3) yields

$$\tilde{\mathcal{P}}(N, s) = \left[\frac{\lambda}{s + \lambda} \right]^N \frac{1}{s + \lambda}. \quad (9.3.4)$$

Inverting the Laplace transform gives the standard Poisson distribution:

$$\mathcal{P}(N, t) = \frac{(\lambda t)^N}{N!} e^{-\lambda t}. \quad (9.3.5)$$

Thus, a renewal process with exponential waiting times reduces to a Poisson process: the number of events in time t follows the Poisson distribution with mean λt . The exponential waiting-time distribution is unique in possessing the memoryless property; all other waiting-time distributions lead to non-Markovian generalizations of the Poisson process.

9.3.2 Incorporating jumps — complete CTRW model

Having constructed a general renewal process for the waiting times, we now incorporate random jump lengths. Let $\rho(x)$ denote the PDF for the displacement of a single jump (the jump-length distribution). We assume the jump-lengths are i.i.d. and independent of the waiting times. Our goal is to compute the PDF $P(x, t)$ of the total displacement x after an observation time t .

Suppose that N jump events occur in $(0, t)$. Let $\mathcal{P}(N, t)$ be the probability for this to occur. If N jumps occur and each jump length is drawn from $\rho(x)$, the distribution of the sum of the N i.i.d. random displacements is the N -fold convolution of ρ with itself. Denote this convolution by $\rho^{\otimes N}(x)$. Then the overall displacement PDF is

$$P(x, t) = \sum_{N=0}^{\infty} \mathcal{P}(N, t) \rho^{\otimes N}(x). \quad (9.3.6)$$

The term $N = 0$ corresponds to no jumps within $(0, t)$, in which case $x = 0$ with probability one, and the distribution reduces to $\delta(x)$. (Here we implicitly assume that at $t = 0$ the particle is located at the origin.)

A useful way to handle the convolutions in (9.3.6) is to apply the Fourier transform in space and the Laplace transform in time. Define

$$\hat{P}(k, t) = \int_{-\infty}^{\infty} e^{-ikx} P(x, t) dx, \quad \hat{\rho}(k) = \int_{-\infty}^{\infty} e^{-ikx} \rho(x) dx.$$

Taking the Fourier transform of (9.3.6) yields

$$\hat{P}(k, t) = \sum_{N=0}^{\infty} \mathcal{P}(N, t) [\hat{\rho}(k)]^N.$$

Taking the Laplace transform of this expression in time and using (9.3.3) gives

$$\hat{\tilde{P}}(k, s) = \sum_{N=0}^{\infty} \left[\tilde{\phi}(s) \right]^N \frac{1 - \tilde{\phi}(s)}{s} [\hat{\rho}(k)]^N. \quad (9.3.7)$$

Since $|\tilde{\phi}(s)| < 1$ and $|\hat{\rho}(k)| < 1$ for non-zero s and k , the geometric series sums to

$$\hat{\hat{P}}(k, s) = \frac{1 - \tilde{\phi}(s)}{s} \frac{1}{1 - \hat{\rho}(k) \tilde{\phi}(s)}. \quad (9.3.8)$$

This compact expression is the *Montroll–Weiss formula* for the Fourier–Laplace transform of the CTRW propagator $P(x, t)$. It is remarkable because it decouples the jump-length distribution and the waiting-time distribution: $\hat{\rho}(k)$ and $\tilde{\phi}(s)$ enter only via the product $\hat{\rho}(k) \tilde{\phi}(s)$.

In general, it is not possible to invert both the Fourier and Laplace transforms in closed form. However, various limits and special cases yield important insights, as we now discuss.

9.4 Brownian motion as an asymptotic limit of CTRW

The continuous time random walk (CTRW) model is much more general than ordinary diffusion: it allows both the lengths of jumps and the waiting times between jumps to have *arbitrary* probability distributions. Remarkably, Brownian motion emerges from this discrete, renewal-process framework in the limit of large observation times and distances. To see why, it is useful to analyse the CTRW propagator in Fourier–Laplace space and then perform asymptotic expansions.

We show that under certain conditions CTRW reduces to Brownian motion in the long-time, long-distance limit. Assume that the jump-length distribution $\rho(x)$ has zero mean and finite variance σ^2 . This implies that its Fourier transform admits an expansion for small k :

$$\hat{\rho}(k) = 1 - \frac{\sigma^2 k^2}{2} + \mathcal{O}(k^4). \quad (9.4.1)$$

Assume also that the waiting-time PDF $\phi(t)$ has a finite mean τ ; its Laplace transform then satisfies

$$\tilde{\phi}(s) = 1 - \tau s + \mathcal{O}(s^2) \quad \text{for } s \rightarrow 0. \quad (9.4.2)$$

Importantly, **no assumption is made about the shape** of these distributions beyond their first two moments. The jump lengths need not be normally distributed; they may have exponential, uniform, or other tails—as long as the variance is finite. Likewise, the waiting times need not be Poissonian; the mean waiting time must simply be finite. The macroscopic behaviour is governed entirely by these moments, not by the detailed form of the PDFs.

Inserting (9.4.1) and (9.4.2) into the Montroll–Weiss formula (9.3.8) yields, for small k and s ,

$$\begin{aligned} \hat{\hat{P}}(k, s) &= \frac{1 - (1 - \tau s)}{s} \frac{1}{1 - (1 - \sigma^2 k^2/2)(1 - \tau s)} \\ &= \frac{\tau}{1 - 1 + \sigma^2 k^2/2 + \tau s} + \text{higher-order terms} \\ &= \frac{1}{s + \frac{\sigma^2}{2\tau} k^2}. \end{aligned}$$

Dropping the higher-order terms and identifying $D = \sigma^2/(2\tau)$, we arrive at

$$\hat{\hat{P}}(k, s) = \frac{1}{s + Dk^2}. \quad (9.4.3)$$

The inverse Laplace transform of $1/(s + Dk^2)$ is $e^{-Dk^2 t}$, so in Fourier space

$$\hat{P}(k, t) = \hat{P}(k, 0) e^{-Dk^2 t}.$$

If the initial condition is localized at the origin, $P(x, 0) = \delta(x)$, then $\hat{P}(k, 0) = 1$. The inverse Fourier transform gives the familiar Gaussian propagator:

$$P(x, t) = \frac{1}{\sqrt{4\pi Dt}} \exp\left(-\frac{x^2}{4Dt}\right).$$

Thus, Brownian motion emerges as the hydrodynamic limit of CTRW when the jump-length PDF has finite variance and the waiting-time PDF has finite mean. In this limit the step size (σ) and average waiting time (τ) define an effective diffusion coefficient $D = \sigma^2/(2\tau)$. The central limit theorem ensures that a large number of small, independent jumps produces a Gaussian displacement distribution.

Physically, a random walker executing a CTRW may have large jumps and variable waiting times at microscopic scales. When the motion is viewed on timescales much longer than τ and on length scales much larger than σ , however, the intricate details of each jump blur out, and the trajectory becomes a smooth, continuous Brownian path. The only memory retained of the microscopic dynamics is the variance of the step length (which sets the diffusive step size) and the average waiting time (which sets the rate at which steps occur). If either the variance diverges (e.g. Lévy flights) or the mean waiting time diverges (e.g. heavy-tailed trapping), the Brownian limit does not exist and anomalous diffusion results.

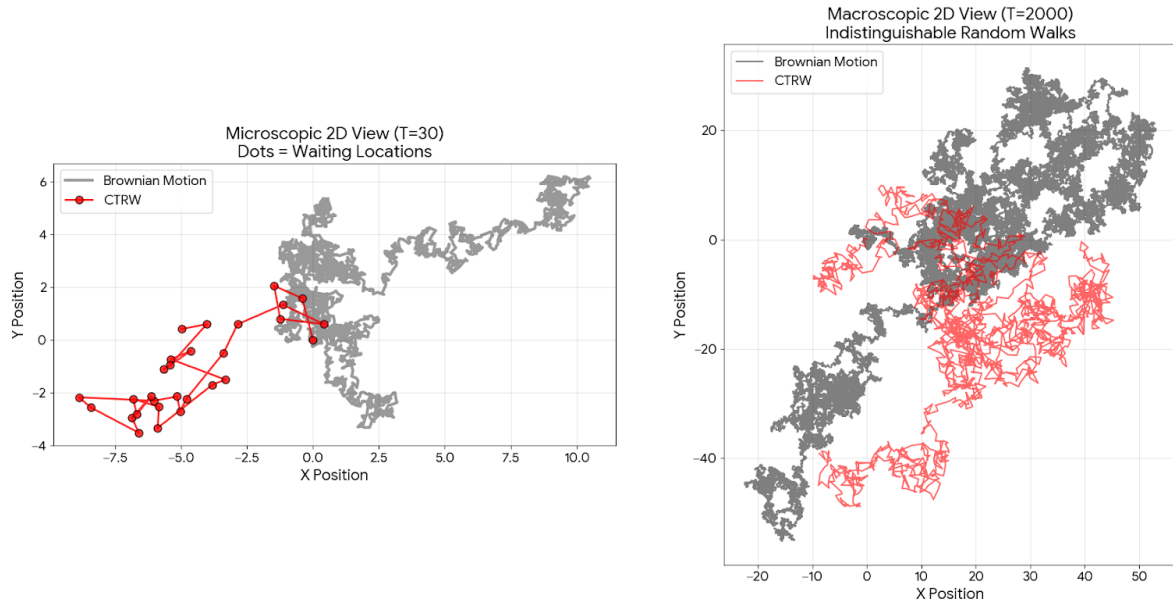


Figure 3: An illustration of the CTRW-to-Brownian limit. The jagged path on the left shows a continuous-time random walk with discrete jumps and waiting times. When the same trajectory is coarse-grained or observed at a much larger scale (right), the details of individual jumps blur together and the path appears as a smooth Brownian motion.

9.5 Water diffusion - description in the CTRW

The CTRW framework encompasses both the Brownian limit and deviations from it at finite time scales and wavevectors. Here we consider the case where the waiting-time distribution is exponential

(a Poisson process) but the jump-length distribution is arbitrary. This is often referred to as a *jump-diffusion* model. We shall show how this model can reproduce or explain the results for water diffusion observed in QENS experiments.

Starting from the Fourier transform of the propagator (time-domain expression) in (9.3.6),

$$\hat{P}(k, t) = \sum_{N=0}^{\infty} \mathcal{P}(N, t) [\hat{\rho}(k)]^N,$$

and substituting the Poisson distribution (9.3.5) for $\mathcal{P}(N, t)$, we obtain

$$\hat{P}(k, t) = \sum_{N=0}^{\infty} \frac{(t/\tau_0)^N}{N!} e^{-t/\tau_0} [\hat{\rho}(k)]^N = e^{-t/\tau_0} \sum_{N=0}^{\infty} \frac{[(t/\tau_0) \hat{\rho}(k)]^N}{N!}.$$

The series is the exponential function, and since $|\hat{\rho}(k)| < 1$ for $k \neq 0$, we get

$$\hat{P}(k, t) = \exp[-(t/\tau_0) (1 - \hat{\rho}(k))] = \exp[-\Gamma(k)t] \quad (9.5.1)$$

where, $\Gamma(k) = (1 - \hat{\rho}(k)) / \tau_0$, is the decay rate at which the correlation function decays. Equation (9.5.1) is the most general form of the characteristic function of a jump-diffusion process with Poisson-distributed waiting times and jump-length distribution $\rho(x)$. The characteristic function, is indeed the self-intermediate scattering function (SISF) discussed in lecture 7. We know that the observed QENS spectra is essentially a time Fourier transform of the SISF, which entails that it will be Lorentzian of the form,

$$S(k, \omega) = \frac{1}{\pi} \frac{\Gamma(k)}{\Gamma(k)^2 + \omega^2} \quad (9.5.2)$$

wherein $\Gamma(k)$ is the half-width at half-maximum of the Lorentzian. We can already note that this will not vary as k^2 , as it will strongly depend on the value of $\hat{\rho}(k)$.

In the case of water diffusion, very specifically, we can consider the jump-length distribution to be a symmetric exponential of the form

$$\rho(x) = \frac{1}{2x_0} \exp[-|x|/x_0] \quad (9.5.3)$$

This is called as the Laplace distribution, it is symmetric about the origin, reflecting the symmetry of the diffusion process. In order to use this to explain the QENS spectra, we need to calculate the Fourier transform of this function,

$$\hat{\rho}(k) = \frac{1}{1 + (kx_0)^2} \quad (9.5.4)$$

Through the relationship $\Gamma(k) = (1 - \hat{\rho}(k)) / \tau_0$, we get,

$$\Gamma(k) = \frac{1}{\tau_0} \frac{(kx_0)^2}{1 + (kx_0)^2} \quad (9.5.5)$$

which reproduces the variation of $\Gamma(k)$ that was observed in the QENS spectra for water (Fig. 2). Typically, $\Gamma(k)$ for small values of k increases as Dk^2 with a diffusion coefficient $D = (x_0^2/\tau_0)$. However, beyond a certain value of k , it doesn't follow k^2 scaling anymore, and eventually saturates at the values $(1/\tau_0)$. This is nicely shown in Fig. 3, which displays $\Gamma(k)$ for different choices of τ_0 . The limit to Brownian motion is recovered in the limit $kx_0 \rightarrow 0$. This behavior nicely reconciles with the observation for the QENS signal from water, which physically reflects that the nature of water diffusion is clearly jump-like, in particular when observed using QENS experiments.




Figure 4: Variation of $\Gamma(k)$ with k^2 for a diffusion process following eq. (9.5.5), shown for different average waiting times τ_0

A general perspective can also be observed for this approach towards the Brownian limit for any choice of jump-distribution in general $\rho(x)$. For small k , expanding $\hat{\rho}(k)$ as in (9.5.4) yields the Brownian limit $\hat{P}(k, t) \approx e^{-t(k^2 x_0^2 / \tau_0)} = e^{-Dk^2 t}$ with $D = x_0^2 / \tau_0$. At finite wavevectors the function $\hat{\rho}(k)$ may deviate significantly from its small- k expansion, leading to non-Gaussian displacement distributions. This feature explains why jump-diffusion models can capture deviations from Gaussian behaviour in experiments such as QENS.

9.6 Slower than diffusion – Subdiffusion

In the Brownian (Poisson) case, the waiting time between successive jumps has a finite mean. Then, after a time t , the walker has typically made a number of steps proportional to t , and the mean-squared displacement (MSD) grows linearly, $\langle x^2(t) \rangle \propto t$. In many physical systems, however, there are *intermittent* dynamics: long periods of immobility are punctuated by rapid jumps. The key mathematical signature is that the waiting time can have *no characteristic scale*, so that its mean diverges. This is the simplest route to **subdiffusion**.

9.6.1 Examples of divergent waiting times (why a mean can fail to exist)

The simplest way to see how an average can diverge is to construct a waiting time distribution in which *rare* events are *so large* that they dominate the mean.

Example 1: A discrete toy model. Let the waiting time take values $\tau_n = 2^n$ with probabilities $p_n = C 2^{-n\alpha}$, where $0 < \alpha < 1$ and C is a normalization constant. Then the mean waiting time is

$$\langle \tau \rangle = \sum_{n=0}^{\infty} \tau_n p_n = \sum_{n=0}^{\infty} 2^n C 2^{-n\alpha} = C \sum_{n=0}^{\infty} 2^{n(1-\alpha)}.$$

Since $1 - \alpha > 0$, this geometric series diverges. The probability of huge waiting times decays, but not fast enough to counteract their size.

Example 2: Continuous heavy tail. Suppose the tail of the waiting-time density behaves as $\phi(\tau) \sim A/\tau^{1+\alpha}$. Then the mean includes an integral of the form

$$\int_0^\infty \tau \phi(\tau) d\tau \sim \int_0^\infty \frac{d\tau}{\tau^\alpha},$$

which diverges for $0 < \alpha < 1$. This is the mechanism we will use in CTRW.

9.6.2 Power-law waiting times and the renewal (counting) process

We now assume the waiting-time PDF has a heavy power-law tail,

$$\phi(\tau) \sim \frac{A}{\tau^{1+\alpha}}, \quad \tau \rightarrow \infty, \quad 0 < \alpha < 1, \quad (9.6.1)$$

with $A > 0$ a constant (the precise short-time form of $\phi(\tau)$ is not important for the long-time scaling). The mean waiting time diverges:

$$\langle \tau \rangle = \int_0^\infty \tau \phi(\tau) d\tau \text{ diverges for } 0 < \alpha < 1.$$

Therefore the process has **no typical waiting time**. This alone is enough to break the Brownian logic “number of steps $\sim t/\tau$ ”.

In CTRW notation, let $N(t)$ be the number of renewal events (jumps) up to time t , and let $\mathcal{P}(N, t)$ be the probability that exactly N events have occurred by time t . From the general renewal construction (derived elsewhere in these notes),

$$\tilde{\mathcal{P}}(N, s) = \left[\tilde{\phi}(s) \right]^N \frac{1 - \tilde{\phi}(s)}{s}. \quad (9.6.2)$$

The long-time behaviour is controlled by the small- s behaviour of $\tilde{\phi}(s)$. For a power-law tail of the form (9.6.1), the Laplace transform has the generic small- s expansion

$$\tilde{\phi}(s) \approx 1 - (s\tau_0)^\alpha + \dots, \quad s \rightarrow 0, \quad (9.6.3)$$

where τ_0 is a characteristic microscopic time scale (set by the short-time regularization of ϕ). The corresponding survival probability is

$$\tilde{\Phi}(s) = \frac{1 - \tilde{\phi}(s)}{s} \approx \frac{(s\tau_0)^\alpha}{s}. \quad (9.6.4)$$

Substituting (9.6.3) into (9.6.2) yields

$$\tilde{\mathcal{P}}(N, s) \approx [1 - (s\tau_0)^\alpha]^N \frac{(s\tau_0)^\alpha}{s}. \quad (9.6.5)$$

9.6.3 Asymptotic growth of the mean number of steps $\langle N(t) \rangle$

The MSD of a CTRW with finite-variance jumps is controlled by how many steps occur up to time t . We therefore compute $\langle N(t) \rangle$ explicitly from the Laplace-space representation.

Define the probability generating function

$$G(z, t) = \sum_{N=0}^{\infty} z^N \mathcal{P}(N, t), \quad |z| \leq 1,$$

and take its Laplace transform using (9.6.2):

$$\begin{aligned}\tilde{G}(z, s) &= \sum_{N=0}^{\infty} z^N \tilde{\mathcal{P}}(N, s) = \sum_{N=0}^{\infty} z^N \left[\tilde{\phi}(s) \right]^N \frac{1 - \tilde{\phi}(s)}{s} \\ &= \frac{1 - \tilde{\phi}(s)}{s} \sum_{N=0}^{\infty} \left[z \tilde{\phi}(s) \right]^N = \frac{1 - \tilde{\phi}(s)}{s} \cdot \frac{1}{1 - z \tilde{\phi}(s)}.\end{aligned}\tag{9.6.6}$$

The mean number of steps follows from

$$\langle N(t) \rangle = \left. \frac{\partial G(z, t)}{\partial z} \right|_{z=1}, \quad \Rightarrow \quad \widetilde{\langle N \rangle}(s) = \left. \frac{\partial \tilde{G}(z, s)}{\partial z} \right|_{z=1}.$$

Differentiate (9.6.6):

$$\frac{\partial \tilde{G}}{\partial z} = \frac{1 - \tilde{\phi}(s)}{s} \cdot \frac{\tilde{\phi}(s)}{(1 - z \tilde{\phi}(s))^2}.$$

Setting $z = 1$ gives the exact Laplace-space identity

$$\widetilde{\langle N \rangle}(s) = \frac{1 - \tilde{\phi}(s)}{s} \cdot \frac{\tilde{\phi}(s)}{(1 - \tilde{\phi}(s))^2} = \frac{\tilde{\phi}(s)}{s(1 - \tilde{\phi}(s))}.\tag{9.6.7}$$

Now insert the small- s form (9.6.3). For $s \rightarrow 0$,

$$1 - \tilde{\phi}(s) \approx (s\tau_0)^\alpha, \quad \tilde{\phi}(s) \approx 1,$$

so (9.6.7) becomes

$$\widetilde{\langle N \rangle}(s) \approx \frac{1}{s(s\tau_0)^\alpha} = \frac{1}{\tau_0^\alpha} s^{-(1+\alpha)}.\tag{9.6.8}$$

We now invert the Laplace transform using the standard rule

$$\mathcal{L}^{-1}\{s^{-(1+\alpha)}\} = \frac{t^\alpha}{\Gamma(1+\alpha)}.$$

Therefore, at long times

$$\langle N(t) \rangle \sim \frac{1}{\Gamma(1+\alpha)} \left(\frac{t}{\tau_0} \right)^\alpha, \quad t \rightarrow \infty.\tag{9.6.9}$$

This sublinear growth of the number of steps is the central mechanism of subdiffusion: the walker does move, but it makes fewer and fewer steps per unit time as t increases, because rare, extremely long waiting times dominate.

9.6.4 Derivation of the subdiffusive mean-squared displacement

We now convert the sublinear step count (9.6.9) into the asymptotic MSD.

Assume the jump lengths $\{\Delta x_i\}$ are IID, symmetric, and have finite variance. Let

$$\langle \Delta x \rangle = 0, \quad \langle (\Delta x)^2 \rangle = \sigma^2 < \infty.$$

After exactly N steps, the displacement is

$$x_N = \sum_{i=1}^N \Delta x_i.$$

Compute the conditional MSD at fixed N :

$$\begin{aligned} \langle x_N^2 \rangle &= \left\langle \left(\sum_{i=1}^N \Delta x_i \right)^2 \right\rangle = \left\langle \sum_{i=1}^N (\Delta x_i)^2 + 2 \sum_{i<j} \Delta x_i \Delta x_j \right\rangle \\ &= \sum_{i=1}^N \langle (\Delta x_i)^2 \rangle + 2 \sum_{i<j} \langle \Delta x_i \rangle \langle \Delta x_j \rangle \quad (\text{independence}) \\ &= N\sigma^2 + 2 \sum_{i<j} 0 = N\sigma^2. \end{aligned} \tag{9.6.10}$$

Now average over the random number of steps $N(t)$:

$$\begin{aligned} \langle x^2(t) \rangle &= \sum_{N=0}^{\infty} \mathcal{P}(N, t) \langle x_N^2 \rangle = \sum_{N=0}^{\infty} \mathcal{P}(N, t) (N\sigma^2) \\ &= \sigma^2 \sum_{N=0}^{\infty} N \mathcal{P}(N, t) = \sigma^2 \langle N(t) \rangle. \end{aligned} \tag{9.6.11}$$

Finally, substitute the asymptotic form (9.6.9):

$$\boxed{\langle x^2(t) \rangle \sim \frac{\sigma^2}{\Gamma(1+\alpha)} \left(\frac{t}{\tau_0} \right)^\alpha, \quad 0 < \alpha < 1, \quad t \rightarrow \infty.} \tag{9.6.12}$$

This is **subdiffusion**: the MSD grows as a power law with exponent $\alpha < 1$, slower than the Brownian result $\langle x^2(t) \rangle \propto t$.

9.6.5 Physical interpretation: subdiffusion in crowded biological environments

The subdiffusive result

$$\langle x^2(t) \rangle \sim t^\alpha, \quad 0 < \alpha < 1,$$

does not imply that particles are intrinsically slow or immobile. Instead, it reflects how *time is spent* during the motion. In CTRW-based subdiffusion, the dominant mechanism is the presence of broadly distributed waiting times between successive displacements. Most steps are separated by short waiting times, but rare waiting periods can be extremely long, and these rare events dominate the long-time statistics.

What the result does *not* assume. It is important to stress what is *not* required to obtain subdiffusion in CTRW:

- The jump lengths need not be anomalous; they may be microscopic and have finite variance.
- There is no need for spatial disorder to be fractal or long-range correlated.
- The particle itself does not need to have memory.

All anomalous behaviour arises from the statistics of the *waiting times*.

Protein motion inside a cell: a physical picture. Inside a living cell, a protein diffuses in a highly crowded, heterogeneous environment: the cytoplasm contains other proteins, nucleic acids, membranes, cytoskeletal filaments, and organelles. From the point of view of a tracer protein, this environment creates a sequence of transient “trapping” events. Physically, these may correspond to:

- transient binding to other macromolecules,
- temporary confinement within a mesh of cytoskeletal filaments,
- caging by surrounding crowding agents,
- local changes in viscoelastic response of the cytoplasm.

Each such event immobilizes the protein for some time before it is released and resumes diffusion. Crucially, the durations of these immobilization periods are not narrowly distributed. Some interactions are weak and short-lived, while others are strong or geometrically constrained and can last orders of magnitude longer.

Why waiting times become heavy-tailed. Escape from a local trap typically requires overcoming an energy barrier or a structural rearrangement. If the distribution of barrier heights is broad, the associated escape times,

$$\tau \sim \tau_0 \exp(E/k_B T),$$

become broadly distributed as well. Even a modest spread in barrier energies translates into a power-law tail in the waiting-time distribution. As a result, there is no characteristic time scale for how long a protein remains trapped before moving again.

Connection to subdiffusion. From the CTRW perspective, the protein performs ordinary spatial steps when it moves, but the number of such steps up to time t grows only as t^α . The protein is therefore *mobile*, yet it spends an increasing fraction of time waiting. This leads directly to subdiffusive scaling of the MSD:

$$\langle x^2(t) \rangle \propto \langle N(t) \rangle \sim t^\alpha.$$

Aging and biological implications. An important consequence of heavy-tailed waiting times is aging: the longer the protein has remained trapped, the longer it is expected to remain trapped in the future. This means that time-translation invariance is broken, and correlation functions depend on the history of the system. In biological contexts, this can manifest as:

- time-dependent mobility of proteins,
- non-exponential relaxation in fluorescence or tracking experiments,
- apparent slowdown of transport at long observation times.

Summary. Subdiffusion of proteins in cells does not require exotic spatial structures or anomalous forces. It can arise naturally from intermittent dynamics in a crowded, heterogeneous environment, where trapping times have no characteristic scale. CTRW with power-law waiting times provides a minimal and physically transparent framework to describe this behaviour, linking microscopic trapping events to macroscopic anomalous transport.

9.7 Summary and outlook

In this chapter we have developed a detailed derivation of the continuous time random walk (CTRW) model and related it to both microscopic kinetic theory and macroscopic diffusion. By randomizing the waiting times and jump lengths independently, CTRW provides a flexible framework that can describe both Poissonian and non-Poissonian arrival processes and both Gaussian and non-Gaussian spatial displacements. We showed how the standard Brownian (diffusion) model emerges as the hydrodynamic limit of CTRW when the waiting-time distribution has finite mean and the jump-length distribution has finite variance. We also derived the Montroll–Weiss formula and discussed jump-diffusion models that explain experimental observations such as deviations from hydrodynamic scaling in water self-diffusion. Lastly, we found that sub-diffusion (slower than diffusion) phenomena observed in biological systems can be explained within CTRW framework, when we consider divergent mean-waiting times, i.e. when there are rare events in which particles are trapped for extremely long-times.